arXiv:1610.05207v1 [math.GT] 17 Oct 2016

LINK COBORDISMS AND FUNCTORIALITY IN LINK FLOER HOMOLOGY

IAN ZEMKE

Abstract. We construct cobordism maps associated to decorated link cobordisms for link Floer homology. The maps are associated to oriented surfaces with divides embedded in 4-dimensional cobordisms between 3-manifolds with decorated links. The maps are defined on an equivariant, filtered, curved chain homotopy type invariant, denoted $\mathcal{CFL}$. We give a rigorous proof of invariance, and prove several important properties. We also make a comparison with the graph TQFT defined by the author for Heegaard Floer homology, showing that the maps associated to cobordisms with embedded, cyclically ordered graphs can be recovered from the link Floer TQFT, giving an alternate construction and proof of invariance of those maps.

Contents

1. Introduction 2
1.1. Comparison of curved link Floer homology to other versions of link Floer homology 6
1.2. Comparison to other link cobordism map constructions 6
1.3. Organization 7
1.4. Acknowledgments 7
2. Decorated link cobordisms, the complexes $\mathcal{CFL}^\circ(\mathcal{H},\mathfrak{s})$ and the coherent chain homotopy type invariant $\mathcal{CFL}^\circ(Y,\mathbb{L},\mathfrak{s})$ 7
2.1. The category of decorated link cobordisms 7
2.2. Curved chain complexes 8
2.3. Heegaard diagrams for multi-based links 9
2.4. Spin^c structures and admissibility 9
2.5. The chain complexes $\mathcal{CFL}^\circ(\mathcal{H},\mathfrak{s})$ 9
2.6. Colorings of complexes, and the category of curved $\mathbb{Z}^{\mathfrak{P}}$-filtered, $\mathbb{Z}_2[U_{\mathfrak{P}}]$-chain complexes 10
3. Background on $\mathcal{CFL}^\circ$, quasi-stabilization, basepoint maps, and disk stabilization 11
3.1. The modules $\mathcal{HFL}^-$ for unlinks in $(S^1\times S^2)^{\#g}$ 11
3.2. Quasi-stabilization 13
3.3. The maps Φ_w and Ψ_z and relations amongst them 15
3.4. Relations between the quasi-stabilization maps $S^\pm_{w,z}$ and $T^\pm_{w,z}$ and the maps Φ_w and Ψ_z 17
3.5. Quasi-stabilization and basepoint moving maps I 20
3.6. Quasi-stabilization and basepoint moving maps II 23
3.7. Disk-stabilization maps $S^\pm_{\mathbb{U},D}$ 25
4. Maps for handles attached away from L 26
4.1. Model link cobordisms for handles away from L 27
4.2. Maps for framed 0- or 2-spheres 27
4.3. Maps for framed 1-dimensional links 29
5. Quasi-stabilization and triangle maps 30
5.1. Strongly positive diagrams for $(S^1\times S^2)^{\#k}$ and quasi-stabilization of Heegaard triples 30
5.2. Weakly positive diagrams and further holomorphic triangle computations 32
6. Definition of the band maps 34
6.1. Topological preliminaries on bands 35
6.2. Maps for type-**z** bands 36
6.3. Maps for type-**w** bands 38
7. Properties and relations involving the band maps 38
7.1. Band maps and quasi-stabilization 38

7.2. Band maps and disk stabilization 41
7.3. Commuting of band maps I 44
7.4. Band maps and basepoint moving maps 45
7.5. Commuting band maps II 47
7.6. A final relation between the $F_B^{\mathbf{z}}$ and $F_B^{\mathbf{w}}$ maps. 51
8. Maps for framed 0-spheres along L 52
8.1. Model compound 1-handle/band cobordisms 52
8.2. Construction of the compound 1-handle/band maps 53
8.3. A remark on non-nullhomologous links and Alexander gradings 54
9. Parametrized Kirby decompositions of link cobordisms 55
9.1. Parametrized cobordisms and elementary parametrized cobordisms 55
10. Morse theory and moves between parametrized Kirby decompositions 57
10.1. Weak equivalences of parametrized Kirby decompositions 58
10.2. Spaces of adapted coordinates 61
10.3. Changing the Morse function away from Σ 64
10.4. Changing the Morse function f_Σ on Σ 65
10.5. Morse data and parametrized Kirby decompositions 69
10.6. Proof of Proposition 10.3 71
11. Construction and invariance of the cobordism maps 72
11.1. Maps induced by elementary parametrized link cobordisms of type (1) and (2) 73
11.2. Maps induced by elementary parametrized link cobordisms of type (3) 74
11.3. Maps induced by elementary parametrized link cobordisms of type (4) 75
11.4. Definition and invariance of the cobordism maps 77
12. Maps for cobordisms where $\pi_0(\mathbb{L}_i) \to \pi_0(\Sigma)$ is not surjective 81
13. The composition law 84
14. Link functors and graph TQFTs 85
14.1. The ambiguity of $(\Sigma(\Gamma), \mathcal{A}(\Gamma))$ 87
14.2. The maps Ψ_z, A_λ and the proof of Theorem C 87
References 89

1. Introduction

In [OS04a] and [OS06], Ozsváth and Szabó construct algebraic invariants of 3-manifolds, called Heegaard Floer homology, such that cobordisms between 3-manifolds induce maps between the Heegaard Floer homologies of the ends. Defined by Ozsváth and Szabó in [OS04b] and independently by Rasmussen in [Ras03], there is a refinement of Heegaard Floer homology for knots embedded in 3-manifolds, called knot Floer homology. In [OS08], Ozsváth and Szabó define a generalization called link Floer homology, for links embedded in three-manifolds.

To a four manifold W with properly embedded surface $\Sigma \subseteq W$, the ends $\partial\Sigma$ are links in ∂W. In light of the cobordism maps in [OS06] associated to 4-manifolds, it is natural to expect functorial maps between the link Floer homologies of the ends of the link cobordism (W, Σ). In this paper we provide such a construction of a $(3+1)$-dimensional TQFT for a version of the full link Floer complex, and prove invariance of the maps.

We call the version of link Floer homology we will use **curved link Floer homology**. A multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link embedded in a 3-manifold with basepoints $\mathbf{w}$ and $\mathbf{z}$, such that each component has at least two basepoints, and the basepoints alternate between $\mathbf{w}$ and $\mathbf{z}$ as one travels along the link. To an oriented multi-based link $\mathbb{L}$ in a 3-manifold Y with a Spin^c structure $\mathfrak{s} \in \text{Spin}^c(Y)$, we will define algebraic objects

$$\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}).$$

In this paper we will focus on the versions $\circ \in \{-, \infty\}$. If $S = \{s_1, \dots, s_n\}$ is a set, we define $\mathbb{Z}_2[U_S]$ to be the polynomial ring generated by the formal variables $U_{s_1}, \dots, U_{s_n}$. To a diagram $\mathcal{H}$ of the link $\mathbb{L}$, we will define a $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$-module $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ with a filtration of $\mathbb{Z}^{\mathbf{w}} \oplus \mathbb{Z}^{\mathbf{z}}$. The module $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ also has an

endomorphism ∂ which satisfies

$$\partial^2 = k \cdot \mathrm{id},$$

for some $k \in \mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$, depending on the link $\mathbb{L}$. The precise constant k can be found in Lemma 2.9, and is determined by the configuration of basepoints on the link. The fact that k can be nonzero results from the bubbling of Maslov index two holomorphic disks when computing ∂^2. Indeed the fact that such bubbling occurs should be thought of as an essential feature of the link Floer package, as it leads to interesting relations between the maps which appear.

Modules with an endomorphism squaring to the action of a scalar in the ground ring are often referred to as "matrix factorizations", and appear in the context of Khovanov-Rozansky homology. See [KR08]. We will call a module with such an endomorphism a **curved chain complex**.

Thus for each diagram $\mathcal{H}$, we get a curved, filtered chain complex $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ over $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$. The filtered, $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$-equivariant, curved chain homotopy type is an invariant of $(Y, \mathbb{L}, \mathfrak{s})$. To define maps associated to link cobordisms, we need to consider a more subtle invariant, involving canonical homotopy equivalences between the complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ for various diagrams $\mathcal{H}$. To this end, we make the following definition:

Definition 1.1. *If $\mathcal{C}$ is an category and A is a set, a **transitive system over $\mathcal{C}$ indexed by** A is a map $C : A \to \mathcal{C}$, together with distinguished morphisms*

$$\Phi_{a \to a'} : C(a) \to C(a')$$

for $a, a' \in A$ such that

- $\Phi_{a \to a} = \mathrm{id}$;
- $\Phi_{a' \to a''} \circ \Phi_{a \to a'} = \Phi_{a \to a''}$.

*If C_1 and C_2 are two transitive systems over $\mathcal{C}$, indexed by A_1 and A_2 respectively, then a **morphism of transitive systems** over $\mathcal{C}$ is a collection of morphisms F_{a_1,a_2}, indexed by pairs $(a_1, a_2) \in A_1 \times A_2$, such that*

$$F_{a_1,a_2} : C_1(a_1) \to C_2(a_2)$$

and the following diagram commutes for any two pairs $(a_1, a_2), (a_1', a_2') \in A_1 \times A_2$:

$$\begin{array}{ccc} C_1(a_1) & \xrightarrow{F_{a_1,a_2}} & C_2(a_2) \\ \downarrow{\scriptstyle \Phi_{a_1 \to a_2}} & & \downarrow{\scriptstyle \Phi_{a_2 \to a_2'}} \\ C_1(a_1') & \xrightarrow{F_{a_1',a_2'}} & C_2(a_2') \end{array}.$$

The complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ fit into the structure of a transitive system over $\mathcal{C}$ indexed by A, where A is the set of all admissible diagrams for $(Y, \mathbb{L})$, and $\mathcal{C}$ is the category of chain $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$-equivariant chain complexes with a filtration of $\mathbb{Z}^{\mathbf{w}} \oplus \mathbb{Z}^{\mathbf{z}}$. The morphisms of this category of $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$-equivariant homomorphisms modulo filtered equivariant chain homotopies. To a pair of diagrams, $\mathcal{H}$ and $\mathcal{H}'$ for a link, we can define a map

$$\Phi_{\mathcal{H} \to \mathcal{H}'} : \mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\mathcal{H}', \mathfrak{s}),$$

which is specified up to filtered, $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$-equivariant chain homotopy using the results of [JT12]. We define $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ to be this transitive system over $\mathcal{C}$ and call it the **coherent, filtered, equivariant chain homotopy type**.

If $\mathcal{C}$ is a category of chain complexes, to a transitive system over $\mathcal{C}$, we can associate a single group by first taking the homology of each group, then taking the inverse limit using the isomorphisms $\Phi_{a \to a'}$. The homology groups of link Floer homology do not contain as much information as the chain complexes (and furthermore we can't take homology if the differential doesn't square to zero), so this isn't as useful for our purposes. As the chain homotopies $\Phi_{\mathcal{H} \to \mathcal{H}'}$ are chain homotopy equivalences and not chain isomorphisms, and are only defined up to filtered chain homotopy equivalence, it is unclear how to define a single curved chain complex using an inverse limit construction. As such, we will work with the notion of transitive systems indexed by a set. We note that a morphism between transitive systems is determined by a single one of the maps F_{a_1,a_2}.

To define cobordism maps, we need to work over a fixed category $\mathcal{C}$, while the objects $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ are transitive systems over a category $\mathcal{C}$ which depends on Y and $\mathbb{L}$. To deal with this, we must work with

colored complexes. A coloring of a link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a pair $(\sigma, \mathfrak{P})$ where $\mathfrak{P}$ is a collection of formal variables and $\sigma : \mathbf{w} \cup \mathbf{z} \to \mathfrak{P}$ is a map assigning each basepoint a formal variable.

We form the colored complexes

$$\mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$$

by tensoring with a certain module to identify variables of the same color. The objects $\mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$ are now transitive systems over $\mathcal{C}_{\mathfrak{P}}$, where $\mathcal{C}_{\mathfrak{P}}$ is the category of $\mathbb{Z}_2[U_{\mathfrak{P}}]$-equivariant, $\mathbb{Z}^{\mathfrak{P}}$-filtered chain complexes, and morphisms are chain maps up to to $\mathbb{Z}_2[U^{\mathfrak{P}}]$-equivariant, $\mathbb{Z}^{\mathfrak{P}}$-filtered chain homotopy.

A **decorated link cobordism**

$$(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$$

is a pair where W is a 4-manifold with $\partial W = -Y_1 \sqcup Y_2$, and F is a tuple $F = (\Sigma, \mathcal{A}, \sigma)$ such that

- $\Sigma \subseteq W$ is a properly embedded, oriented surface with $\partial \Sigma = -L_1 \sqcup L_2$;[1]
- $\mathcal{A}$ is a properly embedded 1-manifold in Σ such that the components of $\Sigma \setminus \mathcal{A}$ are partitioned into two subsurfaces, $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, such that $\Sigma_{\mathbf{w}}$ contains all of the $\mathbf{w}$-basepoints on the ends, and $\Sigma_{\mathbf{z}}$ contains all of the $\mathbf{z}$-basepoints on the ends. Further, all curves in $\mathcal{A}$ are part of the boundary of both $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$;
- $\sigma : C(\Sigma \setminus \mathcal{A}) \to \mathfrak{P}$ is a coloring of the components of $\Sigma \setminus \mathcal{A}$.

We call $F = (\Sigma, \mathcal{A}, \sigma)$ a **colored surface with divides**. We include a picture of the kind of cobordism we are considering in Figure 1.1.

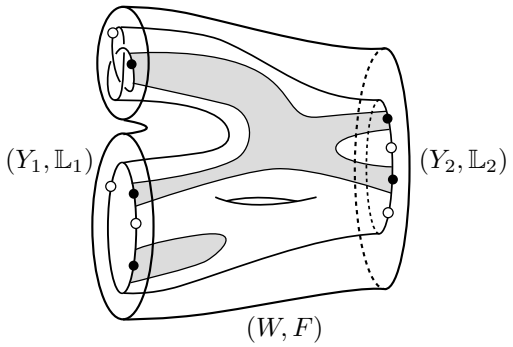

FIGURE 1.1. An example of a colored link cobordism (W, F). The black dots are $\mathbf{w}$-basepoints and the white dots are the $\mathbf{z}$-basepoints. The shaded region is of type-$\mathbf{w}$, and the unshaded region is of type-$\mathbf{z}$.

Our main result is the following:

Theorem A. *To an oriented, decorated link cobordism $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$, where F is a surface with divides with coloring $(\sigma, \mathfrak{P})$, for $\mathfrak{s} \in \mathrm{Spin}^c(W)$, there are maps*

$$F_{W,F,\mathfrak{s}} : \mathcal{C}FL^\circ(Y_1, \mathbb{L}_1, \sigma|_{\mathbb{L}_1}, \mathfrak{P}, \mathfrak{s}|_{Y_1}) \to \mathcal{C}FL^\circ(Y_2, \mathbb{L}_2, \sigma|_{\mathbb{L}_2}, \mathfrak{P}, \mathfrak{s}|_{Y_2}).$$

They are invariants, up to $\mathbb{Z}^{\mathfrak{P}}$-filtered, $\mathbb{Z}_2[U_{\mathfrak{P}}]$-equivariant chain homotopy.

In [OS04a] and [OS06], Heegaard Floer homology is defined only for connected 3-manifolds, and maps are only defined for connected 4-manifolds. The essential limitation in that context was the definition of 1- and 3-handles. For the hat flavor, this is resolved in [Juh15a], but for HF^-, HF^+ and HF^∞, this is resolved in [Zem15b], and $CF^\circ(Y, \mathbf{w}, \mathfrak{s})$ is defined for possibly disconnected 3-manifolds (as long as there is a basepoint in each component), and maps are defined for possibly disconnected 4-manifolds (at the expense of keeping track of a graph or collection of paths which connects the basepoints in the ends). In our present context, we make the following remark:

[1]Orientation conventions will be described in Section 2.

Remark 1.2. The objects $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ can be defined as long as $\pi_0(\mathbb{L})\to\pi_0(Y)$ is a surjection. Cobordism maps for a link cobordisms (W,F) are well defined as long as $\pi_0(\partial\Sigma)\to\pi_0(\partial W)$, $\pi_0(\Sigma)\to\pi_0(W)$, and $\pi_0(\mathcal{A})\to\pi_0(\Sigma)$ are all surjections (note that we don't even require that each component of W have an incoming end or an outgoing end).

As one would expect, the standard Spin^c composition law is also satisfied:

Theorem B. *Suppose that* $(W,F):(Y_1,\mathbb{L}_1)\to(Y_2,\mathbb{L}_2)$ *is an arbitrary link cobordism and*

$$(W,F)=(W_2,F_2)\circ(W_1,F_1)$$

where each (W_i,F_i) *is a link cobordism. If* $\mathfrak{s}_1$ *and* $\mathfrak{s}_2$ *are* Spin^c *structures on* W_1 *and* W_2 *respectively, then we have*

$$F_{W_2,F_2,\mathfrak{s}_2}\circ F_{W_1,F_1\mathfrak{s}_1}\simeq\sum_{\substack{\mathfrak{s}\in\mathrm{Spin}^c(W)\\ \mathfrak{s}|_{W_i}=\mathfrak{s}_i}}F_{W,F,\mathfrak{s}}.$$

The construction of the maps $F_{W,F,\mathfrak{s}}$ goes by way of decomposing the pair (W,F) into elementary pieces, and then defining the maps for each piece. We will define maps corresponding to 1-,2- and 3-handle attachment away from the link, as well as maps for cylindrical cobordisms with simple dividing sets, as well as maps for attaching a band to a link in a 3-manifold. In Figure 1.2, we summarize the maps associated to several simple dividing sets which can appear on cylindrical surfaces $L\times[0,1]$ inside of a copy of $Y\times[0,1]$.

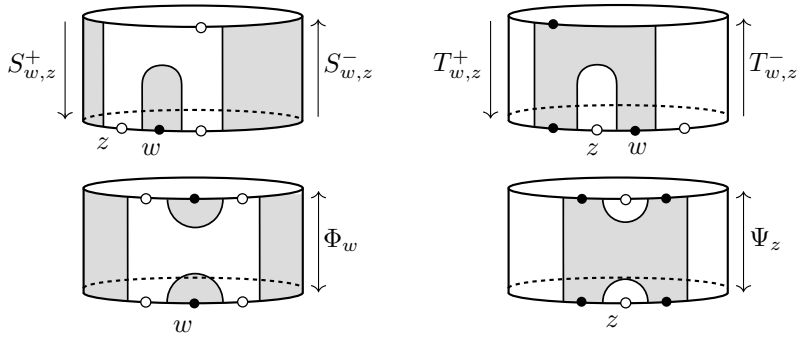

FIGURE 1.2. Several simple dividing sets on cylinders and the maps they correspond to. These are surfaces of the form $\Sigma=L\times[0,1]$ inside of cobordisms $W=Y\times[0,1]$.

The maps on $\mathcal{C}FL^\circ$ for the surfaces appearing in Figure 1.2 appear (using slightly different notation) in [Zem16], where they are used to prove a formula for the map induced by moving all of the basepoints in a loop around a link component. We now briefly discuss the interplay between the algebraic properties of the maps and the colorings of the surface. The map Ψ_z is an endomorphism on $\mathcal{C}FL^\circ$ which satisfies

$$\partial\Psi_z+\Psi_z\partial=U_w+U_{w'},$$

where w and w' are basepoints adjacent to z. If we set $U_w=U_{w'}$ by coloring, then Ψ_z becomes a chain map. Likewise if we look at the picture in Figure 1.2 for the cylinder corresponding Ψ_z, we see that in order to color the surface corresponding to this diagram, the basepoints w and w' must have the same color. The interplay between the algebraic requirements for maps to be well defined on $\mathcal{C}FL^\circ$ and the coloring of surfaces will be a common thread through this paper.

Additionally, we define maps associated to oriented bands attached to links in a 3-manifold. A generic Morse function on Σ will have critical points which do not intersect the divides $\mathcal{A}$. As such, a critical point can occur in either a type-$\mathbf{w}$ region or a type-$\mathbf{z}$ region. Similarly, it turns out that there are two natural ways to define band maps on the modules $\mathcal{C}FL^\circ$. As such, we define two band maps, denoted $F^{\mathbf{w}}_B$ and $F^{\mathbf{z}}_B$, depending on whether the band is to take the designation of type-$\mathbf{w}$ or type-$\mathbf{z}$ as a part of the surface with divides.

By decomposing a link cobordism into pieces, we can define link cobordism maps as a composition involving the maps $S^\pm_{w,z}$ and $T^\pm_{w,z}$ for adding and removing basepoints, the two band maps, and handle attachment

maps for attaching four dimensional handles away from the link L. We then develop the Morse theory necessary to prove that the maps associated to a link cobordism do not depend on the decomposition.

Finally, we will also compare the maps constructed in this paper for $\mathcal{C}FL^\circ$ to the cobordism maps constructed in [Zem15b] for cobordisms (W,Γ) with embedded graphs between closed, multi-pointed 3-manifold. A graph cobordism $(W,\Gamma) : (Y_1,\mathbf{w}_1) \to (Y_2,\mathbf{w}_2)$ is a pair such that W is a compact 4-manifold with $\partial W = -Y_1 \sqcup Y_2$, and Γ is a graph, properly embedded in W, such that $\Gamma \cap Y_i = \mathbf{w}_i$, and such that the basepoints in both Y_1 and Y_2 have valence one in Γ. A ribbon graph Γ is a graph with cyclic orderings at each vertex. A colored graph Γ is a graph with an assignment of a formal variable to each component of Γ. In [Zem15b], the author proves that there are Spin^c-functorial maps $F_{W,\Gamma,\mathfrak{s}}$ associated to colored ribbon graph cobordisms (W,Γ) with a 4-dimensional Spin^c structure.

The requirement in the graph TQFT from [Zem15b] that graphs possess cyclic orderings of the vertices seems very suggestive of a relationship between the graph TQFT and a TQFT involving surfaces embedded in a 4-manifold. Indeed from a ribbon graph cobordism, one can construct a link cobordism

$$(W, F(\Gamma)) : (Y_1, \mathbb{U}_1) \to (Y_2, \mathbb{U}_2),$$

where $\mathbb{U}_i$ is a collection of unknots in Y_i, such that each component of $\mathbb{U}_i$ contains a single basepoint of $\mathbf{w}_i$. The surface with divides $F(\Gamma)$ are specified up to a minor ambiguity (see Section 14 for details). If we tensor the $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$-module $\mathcal{C}FL^-(Y_1,\mathbb{U}_1,\mathfrak{s}_1)$ with $\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}}]/I_{V=1}$, where $I_{V=1}$ denotes the ideal generated by polynomials of the form $(V_z - 1)$, for $z \in \mathbf{z}$, then we recover

$$CF^-(Y_1,\mathbf{w}_1,\mathfrak{s}_1).$$

A similar statement holds for the other flavors of CF°. The maps associated to the link cobordism $(W,F(\Gamma))$ thus yield maps

$$\widetilde{F}_{W,\Gamma,\mathfrak{s}} : CF^-(Y_1,\mathbf{w}_1,\sigma_1,\mathfrak{P},\mathfrak{s}_1) \to CF^-(Y_2,\mathbf{w}_2,\sigma_2,\mathfrak{P},\mathfrak{s}_2).$$

We prove the following:

Theorem C. *Given a colored ribbon graph cobordism (W,Γ), we let $F(\Gamma)$ denote a colored, oriented surface with divides, constructed as in Section 14. After setting the $V_{\mathbf{z}}$ variables equal to 1, the maps $\widetilde{F}_{W,\Gamma,\mathfrak{s}}$, defined as the link cobordism maps $F_{W,F(\Gamma),\mathfrak{s}}$, agree up to $\mathbb{Z}_2[U_{\mathfrak{P}}]$-equivariant chain homotopy with the maps $F_{W,\Gamma,\mathfrak{s}}$ defined in* [Zem15b].

1.1. **Comparison of curved link Floer homology to other versions of link Floer homology.** In [OS04b] and [OS08], Ozsváth and Szabó define knot and link Floer homologies for knots and links in rational homology spheres. Curved link Floer homology recovers these theories. For example, if K is a knot in an integer homology sphere, then the object $CFK^\infty(Y,K)$, defined by Ozsváth and Szabó, is a $\mathbb{Z}_2$ subcomplex of $\mathcal{C}FL^\infty(Y,K,\mathfrak{s}_0)$, where $\mathfrak{s}_0$ is the unique element of Spin$^c(Y)$. In fact, in [Zema], we will define an absolute Alexander grading on $\mathcal{C}FL^\infty(Y,K,\mathfrak{s}_0)$ and $CFK^\infty(Y,K)$ can naturally be identified with $\mathcal{C}FL^\infty(Y,K,\mathfrak{s}_0)_0$, the subcomplex in zero Alexander grading. The version of the full link Floer complex from [OS08] can be recovered similarly, by first tensoring $\mathcal{C}FL^\circ$ with a module to identify all of the $V_{\mathbf{z}}$ variables on each link component, and then taking a subcomplex in a single Alexander grading.

One of the advantages of using curved link Floer homology is that it allows us to define maps without paying attention to Alexander gradings or relative Spin^c structures. For example, the maps Φ_w and Ψ_z are not endomorphisms of $CFK^\infty(Y,K)$, as they change the relative Spin^c structure (though the product $\Phi_w\Psi_z$ is an endomorphism). If we were very careful about gradings, we could work with the uncurved complexes, though this would be extremely tedious. Additionally, we will need to work with links which may have components which are not nullhomologous, even rationally. To deal with this, one would need to adapt the more standard $CFL^\infty(Y,L)$ complexes to deal with this.

1.2. **Comparison to other link cobordism map constructions.** There are several examples of constructions like ours already in the literature. The most similar is the one in [Juh15a], where Juhász constructs a TQFT for $\widehat{HFL}$ associated to cobordisms with embedded surfaces, decorated with a dividing set. The resulting maps arising in the cases of knot concordances are studied in [JM15]. It would be interesting to compare our maps, to those defined by Juhász. Indeed our maps do yield maps on $\widehat{HFL}$ by simply setting all of the variables to zero. In such a case, the colorings become irrelevant, and our definition of a decorated

link cobordism becomes essentially equivalent to the one in [Juh15a]. To this end, we make the following conjecture:

Conjecture 1.3. *The maps* $\widehat{F}_{W,F} = \sum_{\mathfrak{s}\in \mathrm{Spin}^c(W)} \widehat{F}_{W,F,\mathfrak{s}}$ *defined in this paper for the hat flavor coincide with the maps* $F_{W,F}$ *defined in* [Juh15a] *using sutured Floer homology.*

In [JM16], Juhász and Marengon provide a full description of all of the link cobordism maps arising from the sutured Floer TQFT, as maps which defined in terms of generators of the complexes, as well as maps defined by counting holomorphic curves on Heegaard triple diagrams. In principle, verification of the above conjecture should be tractable using their description, but we will not pursue it in this paper.

There are additional partial constructions of cobordism maps. For example, in [Sar11b], Sarkar defines maps associated to surfaces in $S^3 \times [0,1]$ in order to provide an alternate proof of the Milnor conjecture using grid diagrams. In [OSS15], Ozsváth, Stipsicz and Szabó use grid diagrams to analyze the unoriented four-ball genus of knots and links. Additionally, there is work of Alishahi and Eftekhary to describe cobordism maps on a minus version of sutured Floer homology in [AE].

Despite these constructions, there is still plenty of room for a new theory, since most constructions either are defined using grid diagrams, admit no straightforward generalization to the full knot Floer complex, or are limited in terms of what 4-manifolds and surfaces they can be applied to. The grid diagram approaches feature the advantage of being defined in a relatively combinatorial fashion, though proving invariance is harder in the setting of grid diagrams, and there is no obvious generalization to links in arbitrary three-manifolds.

1.3. **Organization.** In Section 2 we review the definition of the chain complexes $\mathcal{C}FL^\circ$ and the category of decorated link cobordisms. In Section 3 we describe the maps Φ_w, Ψ_z, the quasi-stabilization maps $S^\pm_{w,z}$ and $T^\pm_{w,z}$, and the maps $S^\pm_{\mathbb{U},D}$, which appear in our construction of the link Floer TQFT. We then prove several important properties and relations of these maps. In Section 4 we describe handle attachment maps for handles which are attached along framed spheres which don't intersect L. In Section 5 we perform some technical triangle map computations. In Section 6 we define maps associated to bands, and prove some important relations. In Section 8 we define cobordism maps associated to a 4-dimensional 1-handle with a 2-dimensional band inside of it. In Section 9, we define the notions of a parametrized cobordism and a parametrized decomposition of a link cobordism. In Section 10 we prove that we can move between any two parametrized decompositions via a set of moves. In Section 11 we give a precise definition of the cobordism maps involved, in terms of maps on Heegaard Floer homology and parametrized decompositions. In Section 12, we provide the construction when the cobordism has a component with an empty end. In Section 13, we prove the Spin^c composition law. In Section 14 we show how the construction here recovers the graph TQFT defined in [Zem15b], and we explain the connection.

1.4. **Acknowledgments.** The author would like to thank Akram Alishahi, András Juhász, Robert Lipshitz, Ciprian Manolescu, and Marco Marengon for helpful conversations and suggestions. The author would also like to thank Sucharit Sarkar and Andy Manion for suggesting to work in the setting of curved chain complexes.

2. Decorated link cobordisms, the complexes $\mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s})$ and the coherent chain homotopy type invariant $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$

In this section, we describe background about the category of decorated link cobordisms, and the definitions of the objects $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ which will appear in this paper.

2.1. **The category of decorated link cobordisms.** In this section we provide the definitions of multi-based links in 3-manifolds, and link cobordisms of colored surfaces with divides between them.

Definition 2.1. *A* ***three manifold with multi-based link*** *is a pair* $(Y,\mathbb{L})$ *where* Y *is an oriented three-manifold and* $\mathbb{L} = (L,\mathbf{w},\mathbf{z})$ *is a tuple where* L *is an oriented link in* Y*, and* $\mathbf{w}$ *and* $\mathbf{z}$ *are collections of basepoints on* L*, such that each component of* L *has at least two basepoints, and the basepoints alternate between those in* $\mathbf{w}$ *and those in* $\mathbf{z}$ *as one traverses the link.*

Definition 2.2. *A **surface with divides** is a pair $(\Sigma, \mathcal{A})$ where Σ is a compact, oriented[2] surface together with a collection $\mathcal{A}$ of properly embedded 1-manifolds such that $\Sigma \setminus \mathcal{A}$ is partitioned into subsurfaces, $\Sigma_{\mathbf{w}}$ and $\Sigma_{\mathbf{z}}$, such that $\Sigma_{\mathbf{w}}$ is on one side of $\mathcal{A}$ and $\Sigma_{\mathbf{z}}$ is on the other.*

To achieve well defined, functorial maps in our link TQFT, we need to consider colorings of the basepoints and the surfaces with divides.

Definition 2.3. *A coloring of a multibased link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a pair $(\sigma, \mathfrak{P})$ where $\mathfrak{P}$ is a finite set of colors and $\sigma : (\mathbf{w} \cup \mathbf{z}) \to \mathfrak{P}$ is an assignment of a color to each basepoint. A coloring of a surface with divides $(\Sigma, \mathcal{A})$ is a pair $(\sigma, \mathfrak{P})$ where $\sigma : C(\Sigma \setminus \mathcal{A}) \to \mathfrak{P}$ is an assignment of a color to each component of $\Sigma \setminus \mathcal{A}$.*

We now define a notion of morphism between three-manifolds with multi-based links.

Definition 2.4. *A pair (W, F) is called a **decorated link cobordism** from $(Y_1, (L_1, \mathbf{w}_1, \mathbf{z}_1))$ to $(Y_2, (L_2, \mathbf{w}_2, \mathbf{z}_2)$ if W is a compact 4-manifold with $\partial W = -Y_1 \sqcup Y_2$, and $F = (\Sigma, \mathcal{A}, \sigma)$ is a colored surface with divides, $\Sigma \subseteq W$ is a properly embedded, compact, oriented surface with $\Sigma \cap W = \partial L_i$. Furthermore, each component of $L_i \setminus \mathcal{A}$ contains exactly one basepoint of $\mathbf{w}_i \cup \mathbf{z}_i$, which is of the same type as the component of $\Sigma \setminus \mathcal{A}$ which it is in.*

To properly define a link cobordism category, we need to work with cobordisms which have a parametrization of the ends. More precisely we define a parametrized link cobordism $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ to be a link cobordism $(W, F) : (Y_1', \mathbb{L}_1') \to (Y_2', \mathbb{L}_2')$ in the previous sense, together with a choice of diffeomorphism $\phi_i : (Y_i', \mathbb{L}_i') \to (Y_i, \mathbb{L}_i)$. Such a definition is necessary to define a proper category, since otherwise the cobordism category lacks identity morphisms, for example. We should always remember that this is the category we are working in, though it's usually cumbersome to make this explicit.

Remark 2.5. One could define a category where link cobordisms with nonorientable surfaces were allowed. It may be possible to define cobordism maps for nonorientable surfaces on a version of $\mathcal{C}FL^\infty$ obtained by identifying the $U_{\mathbf{w}}$ and $V_{\mathbf{z}}$ variables, as in [OSS15], but this is work for a future paper. In this paper we restrict to oriented surfaces.

2.2. **Curved chain complexes.** For the purposes of this section, assume that $\mathcal{R}$ is a ring of characteristic two. In this section we describe the notion of a curved chain complex, which is basically a chain complex where the differential squares to a scalar in $\mathcal{R}$. We could avoid using these if we were willing to either only have two variables (a U-variable and a V-variable), or we are willing to lose some amount of symmetry between the $\mathbf{w}$-basepoints and the $\mathbf{z}$-basepoints, as we did in [Zem16], where we set all of the $V_{\mathbf{z}}$ variables on a link component equal to each other. The setting of curved chain complexes seems like the most natural, so we will use it.

We now precisely describe some terminology (which is elementary in the case of actual chain complexes):

Definition 2.6. *If $\mathcal{R}$ is a ring, a **curved $\mathcal{R}$-chain complex** is an $\mathcal{R}$-module C, together with an endomorphism $\partial : C \to C$ such that*

$$\partial^2 = k \cdot \mathrm{id},$$

*for some $k \in \mathcal{R}$. We say an $\mathcal{R}$-homomorphism $F : (C, \partial) \to (C', \partial')$ of curved $\mathcal{R}$-chain complexes is a **chain map** if $F\partial = \partial F$. An $\mathcal{R}$-equivariant **chain homotopy** between morphisms $F, G : (C, \partial) \to (C', \partial')$ is an $\mathcal{R}$-homomorphism $H : C \to C'$ such that $F + G = \partial' H + H\partial$.*

Note that being chain homotopic is an equivalence relation. We note that the class of curved $\mathcal{R}$-chain complexes is a category, with morphisms being ($\mathcal{R}$-equivariant) chain maps up to ($\mathcal{R}$-equivariant) chain homotopy.

Definition 2.7. *A $\mathbb{Z}^n$ filtration on a curved $\mathcal{R}$-chain complex (C, ∂) is a collection of subsets $\mathcal{F}_I \subseteq C$ ranging over $I \in \mathbb{Z}^n$ such that $\partial \mathcal{F}_I \subseteq \mathcal{F}_I$ and such that if $I \leq J$ (i.e. each component satisfies that inequality) then $\mathcal{F}_I \subseteq \mathcal{F}_J$.*

[2]Our orientation convention that if N is an outward pointing normal vector and $v \in T(\partial\Sigma)$, then v is positively oriented if the pair (v, N) is an oriented basis for $T\Sigma$. This is merely a convention, though it affects the statement of Theorem C, as well as an Alexander grading formula, to appear in [Zema]. This is the natural orientation to choose, as it gives $L \times [0, 1]$, the product orientation, when viewed as a cobordism from $L \times \{0\}$ to $L \times \{1\}$.

2.3. **Heegaard diagrams for multi-based links.** If $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multi-based link in Y^3, we define a **Heegaard diagram** $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ to be an embedded Heegaard diagram of Y such that

(1) $\Sigma \cap L = \mathbf{w} \cup \mathbf{z}$, and Σ divides Y into two solid handlebodies, U_α and U_β,
(2) Σ is oriented with a normal vector pointing into U_β;
(3) L intersects Σ positively at the $\mathbf{z}$-basepoints and negatively at the $\mathbf{w}$-basepoints;
(4) $L \cap U_\alpha$ is isotopic in U_α, in the complement of compressing disks for the $\boldsymbol{\alpha}$ curves, relative to its endpoints to a collection of arcs in Σ;
(5) $L \cap U_\beta$ is isotopic in U_β, in the complement of compressing disks for the $\boldsymbol{\beta}$ curves, relative to its endpoints to a collection of arcs in Σ.

Alternatively, we can connect choose paths on Σ from the $\mathbf{z}$ basepoints to the $\mathbf{w}$ basepoints in the complement of the $\boldsymbol{\beta}$ curves, and choose paths from the $\mathbf{w}$ basepoints to the $\mathbf{z}$ basepoints in the complement of the $\boldsymbol{\alpha}$ curves. Pushing off these collections of curves into U_β and U_α respectively yields a link which is isotopic to L.

To achieve functorial maps, we must work with embedded Heegaard diagrams, i.e. Heegaard diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ equipped with an embedding

$$(\Sigma, \mathbf{w}, \mathbf{z}) \hookrightarrow (Y, \mathbf{w}, \mathbf{z}).$$

An "abstract" Heegaard diagram (a diagram without a choice of embedding) is enough to describe the isomorphism class of the various Heegaard Floer homology invariants $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$, though to describe functorial maps between the complexes $\mathcal{C}FL^\circ$ for various links and 3-manifolds, one needs canonical change of diagrams maps

$$\Phi_{\mathcal{H}\to\mathcal{H}'} : \mathcal{C}FL(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL(\mathcal{H}', \mathfrak{s}),$$

which are defined up to equivariant, filtered chain homotopy. According to [JT12], to describe such maps it is necessary to work with embedded Heegaard diagrams.

2.4. Spin^c **structures and admissibility.** Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is an ℓ-pointed Heegaard diagram for Y. We can define a map

$$\mathfrak{s}_{\mathbf{w}} : \mathbb{T}_\alpha \cap \mathbb{T}_\beta \to \mathrm{Spin}^c(Y).$$

If $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is an intersection point, then we define $\mathfrak{s}_{\mathbf{w}}(\mathbf{x})$, by picking a Morse function defining $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and removing regular neighborhoods of the flowlines passing through the points in $\mathbf{x}$, as well as the basepoints $\mathbf{w}$. This yields a nonvanishing vector field on the complement of this neighborhood of these flowlines, which will extend over these balls, since the flowlines connected critical points with indices of opposite parity. The map $\mathfrak{s}_{\mathbf{w}}$ depends on the choice of basepoints, $\mathbf{w}$, and indeed picking a different basepoints in the same Heegaard diagram changes the map $\mathfrak{s}_{\mathbf{w}}$. We have the following:

Lemma 2.8. *If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, L, \mathbf{w}, \mathbf{z})$ and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, we have*

$$\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) - \mathfrak{s}_{\mathbf{z}}(\mathbf{x}) = PD[L].$$

Proof. This follows from reasoning identical to the proof of [OS04a, Lem. 2.19]. Alternatively, $\mathrm{Spin}^c(Y)$ is described as the set of homology classes of nonvanishing vector fields on Y. Given a closed path $\gamma \in H_1(Y; \mathbb{Z})$, the action of $\gamma \in H_1(Y; \mathbb{Z})$, can be described by a procedure sometimes referred to as Reeb surgery (cf. [Nic03, Section 3.2]). Given the explicit description from [OS04a] of the vector fields $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$, it is easy to see that they differ by Reeb surgery on the link L. □

Hence if the sum of all the components of a link is trivial in $H_1(Y; \mathbb{Z})$, then it doesn't matter whether we take intersection points with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$ or $\mathfrak{s}_{\mathbf{z}}(\mathbf{x}) = \mathfrak{s}$. For a link which is not nullhomologous in this sense, the resulting modules will in general vary. An instructive example is $Y = (S^1 \times S^2)$ with $K = S^1 \times \{pt\}$.

2.5. **The chain complexes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$.** Given a strongly-$\mathfrak{s}$ admissible Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for a link $\mathbb{L}$ in Y^3, we define

$$\mathcal{C}FL^-(\mathcal{H}, \mathfrak{s})$$

to be the free $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$-module with generators of the form

$$U_{\mathbf{w}}^I V_{\mathbf{z}}^J \cdot \mathbf{x},$$

where I and J are nonnegative multi-indices and $\mathbf{x} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$ is an intersection point with $\mathfrak{s}_{\mathbf{w}}(\mathbf{x}) = \mathfrak{s}$. The module

$$\mathcal{C}FL^\infty(\mathcal{H}, \mathfrak{s})$$

is defined as the module with the same generators, but allowing any multi-indices I and J. The module $\mathcal{C}FL^+(\mathcal{H}, \mathfrak{s})$ can be defined as the quotient

$$\mathcal{C}FL^+(\mathcal{H}, \mathfrak{s}) = \mathcal{C}FL^\infty(\mathcal{H}, \mathfrak{s})/\mathcal{C}FL^-(\mathcal{H}, \mathfrak{s}).$$

Now given a multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, we will need to consider diagrams which satisfy an extra admissibility assumption. The admissibility assumption is called strong $\mathfrak{s}$-admissibility, which is described in [OS04a, Def. 4.10], and ensures necessary finiteness results on counts of flowlines.

Given an almost complex structure J_s on $\Sigma \times [0,1] \times \mathbb{R}$ which is suitably generic in the sense of [Lip06], we define an endomorphism

$$\partial : \mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$$

by the formula

$$\partial(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta} \sum_{\substack{\phi \in \pi_2(\mathbf{x}, \mathbf{y}) \\ \mu(\phi) = 1}} \#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y}.$$

The map ∂ does not square to zero, in general. Instead we have the following:

Lemma 2.9 ([Zem16, Lemma 2.1])**.** *The map $\partial : \mathcal{C}FL^\infty(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL^\infty(\mathcal{H}, \mathfrak{s})$ satisfies*

$$\partial^2 = \sum_{K \in |L|} (U_{w_{K,1}} V_{z_{K,1}} + V_{z_{K,1}} U_{w_{K,2}} + U_{w_{K,2}} V_{z_{K,2}} + \cdots V_{z_{K,n_K}} U_{w_{K,1}}),$$

where $w_{K,1}, z_{K,1}, \ldots, w_{K,n_K}, z_{K,n_K}$ are the basepoints on the link component K, in the order that they appear on K.

Thus $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ is a curved chain complex over the ring $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$.

Given two diagrams, $\mathcal{H}$ and $\mathcal{H}'$, for $(Y, \mathbb{L})$, by composing change of diagrams maps corresponding to elementary Heegaard moves (i.e $(1,2)$-stabilization, isotopy relative L, handleslides and isotopies of the $\boldsymbol{\alpha}$- and $\boldsymbol{\beta}$-curves) we can define maps

$$\Phi_{\mathcal{H} \to \mathcal{H}'} : \mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\mathcal{H}', \mathfrak{s})$$

as a composition of maps corresponding to elementary moves. Since the proof of simple handleswap invariance from [JT12] follows from a holomorphic triangle argument, invariance from simple handleswaps holds for $\mathcal{C}FL^\circ$. It follows from [JT12] that up to equivariant, filtered chain homotopy, these maps are independent of the choice of elementary Heegaard moves connecting $\mathcal{H}$ and $\mathcal{H}'$.

Thus we can define $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ as the transitive system over $\mathcal{C}$ indexed by A (cf. Definition 1.1), where A is the set of all strongly $\mathfrak{s}$-admissible diagrams for $(Y, \mathbb{L})$ and $\mathcal{C}$ is the category of $\mathbb{Z}^{\mathbf{w}} \oplus \mathbb{Z}^{\mathbf{z}}$-filtered, $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}]$-equivariant curved chain complexes. We call $\mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s})$ the **coherent, equivariant, filtered, curved chain homotopy type**, which we will usually abbreviate by "coherent chain homotopy type".

2.6. **Colorings of complexes, and the category of curved $\mathbb{Z}^{\mathfrak{P}}$-filtered, $\mathbb{Z}_2[U_{\mathfrak{P}}]$-chain complexes.** An important algebraic step in constructing our link Floer TQFT is coloring complexes. Recall that a coloring $(\sigma, \mathfrak{P})$ of a link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a finite set $\mathfrak{P}$ indexing a collection of formal variables, as well as a map

$$\sigma : \mathbf{w} \cup \mathbf{z} \to \mathfrak{P}.$$

Since we are working with curved complexes instead of standard chain complexes, we no longer need to make any assumptions on $(\sigma, \mathfrak{P})$ to get a well defined curved chain complex (in [Zem16] we assumed that two $\mathbf{z}$-basepoints on a link component were given the same color if they were on the same link component). Given a link $\mathbb{L}$ and a coloring $(\sigma, \mathfrak{P})$, we define $I_{\sigma, \mathfrak{P}}$ as the ideal of $\mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}, U_{\mathfrak{P}}]$ generated by elements of the form $U_w - U_{\sigma(w)}$ and $V_z - U_{\sigma(z)}$. Then we define the coloring module as

$$C_{\sigma, \mathfrak{P}} = \mathbb{Z}_2[U_{\mathbf{w}}, V_{\mathbf{z}}, U_{\mathfrak{P}}]/I_{\sigma, \mathfrak{P}}.$$

Finally we define the colored complexes

$$\mathcal{C}FL^-(\mathcal{H},\sigma,\mathfrak{P},\mathfrak{s}) = \mathcal{C}FL^-(\mathcal{H},\mathfrak{s}) \otimes_{\mathbb{Z}_2[U_{\mathbf{w}},V_{\mathbf{z}}]} C_{\sigma,\mathfrak{P}},$$

and similarly for the ∞ flavor.

The objects $\mathcal{C}FL^\circ(\mathcal{H},\sigma,\mathfrak{P},\mathfrak{s})$ become transitive systems over $\mathcal{C}$, indexed by A, where $\mathcal{C}$ is the category of $\mathbb{Z}^{\mathfrak{P}}$-filtered (the filtration being induced simply by the powers of the variables), curved $\mathbb{Z}_2[U_{\mathfrak{P}}]$-chain complexes, and A is the set of strongly $\mathfrak{s}$-admissible Heegaard diagrams for $(Y,\mathbb{L})$.

We also note that including colorings is essential for functoriality. Without colorings, we do not have a fixed category to work in. Also, most of the maps we define for decorations on the surface are not chain maps until we color the complexes.

Although we will primarily work with curved chain complexes, we will occasionally have the need to actually take homology. In the case that $(\sigma,\mathfrak{P})$ is a coloring of $\mathbb{L}$ such that $\partial^2 = 0$ on $\mathcal{C}FL^\circ(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$, we will write

$$\mathcal{H}FL^\circ(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$$

for the associated homology group.

3. Background on $\mathcal{C}FL^\circ$, quasi-stabilization, basepoint maps, and disk stabilization

In this section, we include several important computations, as well some basic results about the maps $S^\pm_{w,z}, T^\pm_{w,z}$ and $S^\pm_{\mathbb{U},D}$, which appear in the construction of the link Floer TQFT.

3.1. **The modules $\mathcal{H}FL^-$ for unlinks in $(S^1\times S^2)^{\#g}$.** In this section, we analyze $\mathcal{H}FL^-$ for unlinks in connected sums of $(S^1\times S^2)$ with various configurations of basepoints. The analogous statements hold for disjoint unions of such manifolds, though we phrase them only for connected manifolds, for notational simplicity.

We note that when $c_1(\mathfrak{s})$ is torsion, the objects $\mathcal{C}FL^\circ(\mathcal{H},\mathfrak{s})$ have two relative homological gradings, each analogous to the grading on closed 3-manifold invariants defined by Ozsváth and Szabó. We denote one grading by $\mathrm{gr}_{\mathbf{w}}$ and the other by $\mathrm{gr}_{\mathbf{z}}$. The the $U_{\mathbf{w}}$ variables are -2 graded with respect to $\mathrm{gr}_{\mathbf{w}}$ and the $V_{\mathbf{z}}$ variables are 0-graded with respect to $\mathrm{gr}_{\mathbf{w}}$. Similarly the $V_{\mathbf{z}}$ variables are -2 graded with respect to $\mathrm{gr}_{\mathbf{z}}$ and the $U_{\mathbf{w}}$ variables are 0-graded.

For a disk $\phi\in\pi_2(\mathbf{x},\mathbf{y})$, we define

$$\mathrm{gr}_{\mathbf{w}}(\mathbf{x},\mathbf{y}) = \mu(\phi) - 2\sum_{w\in\mathbf{w}} n_w(\phi),$$

and

$$\mathrm{gr}_{\mathbf{z}}(\mathbf{x},\mathbf{y}) = \mu(\phi) - 2\sum_{z\in\mathbf{z}} n_z(\phi).$$

If $c_1(\mathfrak{s})$ is torsion, then $\mathrm{gr}_{\mathbf{w}}$ is independent of the disk ϕ, using the Chern class formula from [OS04]. Analogously, using Lemma 2.8, if $c_1(\mathfrak{s}-PD[L])$ is torsion, then $\mathrm{gr}_{\mathbf{z}}$ will be independent of the disk ϕ.

For $\mathbf{o}\in\{\mathbf{w},\mathbf{z}\}$ we will define $\mathcal{H}FL^-_{\max,\mathbf{o}}(Y,\mathbb{L},\mathfrak{s})$ to denote the subset of $\mathcal{H}FL^-(Y,\mathbb{L},\mathfrak{s}_0)$ of maximal homological grading with respect to $\mathrm{gr}_{\mathbf{o}}$.

Lemma 3.1. *Suppose $\mathbb{U}$ is an unlink in $(S^1\times S^2)^{\#g}$ with ℓ components, and each component has exactly two basepoints. Then there is an isomorphism*

$$\mathcal{H}FL^-_{\max,\mathbf{w}}((S^1\times S^2)^{\#g},\mathbb{U},\mathfrak{s}_0)\cong\mathbb{Z}_2[V_{\mathbf{z}}].$$

In particular, there is an element which generates the maximally $\mathrm{gr}_{\mathbf{w}}$-graded subset with respect to the $\mathbb{Z}_2[V_{\mathbf{z}}]$-module action. Similarly

$$\mathcal{H}FL^-_{\max,\mathbf{z}}((S^1\times S^2)^{\#g},\mathbb{U},\mathfrak{s}_0)\cong\mathbb{Z}_2[U_{\mathbf{w}}],$$

and hence the latter object has an algebraic generator, as well.

The two generators are equal in $\mathcal{H}FL^-$.

Proof. We will prove the claim about $\mathcal{H}FL^-_{\max,\mathbf{w}}$ since the claim about $\mathcal{H}FL^-_{\max,\mathbf{z}}$ follows identically. We can pick g homologically independent curves $\alpha_1,\dots,\alpha_g$ on Σ_g, add a pair of basepoints w_1, and z_1 in the (single) region of $\Sigma_g\setminus(\alpha_1\cup\dots\cup\alpha_g)$.

Now pick $\ell-1$ other points in $\Sigma\setminus(\alpha_1\cup\dots\cup\alpha_g)$, and then add curves $\alpha_{g+1},\dots,\alpha_{g+\ell-1}$ which each bound a disk containing exactly one of the chosen $\ell-1$ points. Call the totality of these curves $\boldsymbol{\alpha}$ and pick $\boldsymbol{\beta}$ to be small Hamiltonian isotopies of the $\boldsymbol{\alpha}$ curves. Now pick basepoints w_i,z_i for $2\leq i\leq\ell$ in each of the disks bounded by the curves $\alpha_{g+1},\dots,\alpha_{g+\ell-1}$ and their corresponding $\boldsymbol{\beta}$-curves.

In this diagram, there is a canonical top degree intersection point Θ. The top degree subset of $\mathcal{C}FL^-(\Sigma_g,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathbf{z},\mathfrak{s}_0)$ is easily seen to be generated by sums of elements of the form $V_{\mathbf{z}}^J\cdot\Theta$. Now we claim that any sum of elements of the form $V_{\mathbf{z}}^J\cdot\Theta$ is a cycle, and no nontrivial sum of such elements is a boundary.

To see that any sum of elements of the form $V_{\mathbf{z}}^J\cdot\Theta$ is a cycle, one simply analyzes the proof of [OS04a, Lem. 9.1]. Using the relative Maslov grading, one sees that if $\phi\in\pi_2(\Theta,\mathbf{y})$ is a Maslov index 1 disk, then

$$2\sum_{w\in\mathbf{w}}n_w(\phi)\leq 0,$$

hence any such disk admitting holomorphic representatives cannot cross any of the $\mathbf{w}$-basepoints. Since the diagram is so simple, it is easy to see that the disks which don't cross over any $\mathbf{w}$-basepoints come in pairs, and hence cancel when counted by the differential. Thus $\partial\Theta=0$.

To see that no nontrivial sum of terms of the form $V_{\mathbf{z}}^J\cdot\Theta$ is a boundary, we note that any disk $\phi\in\pi_2(\mathbf{y},\Theta)$ satisfies

$$\mu(\phi)-2\sum_{w\in\mathbf{w}}n_w(\phi)\leq 0,$$

by the relative homological grading formula, and hence

$$2\sum_{w\in\mathbf{w}}n_w(\phi)\geq 1,$$

implying that if ϕ has a holomorphic representative which is counted by the differential, it must have nonzero multiplicities on the $\mathbf{w}$-basepoints. In particular ϕ must increase the $U_{\mathbf{w}}$-filtration, and hence all cycles live in the subset of $\mathcal{C}FL^-$ of $U_{\mathbf{w}}$-filtration at least 1, and in particular nothing which is a sum of elements of the form $V_{\mathbf{z}}^J\cdot\Theta$ is a cycle. Hence

$$\mathcal{H}FL^-_{\max}((S^1\times S^2)^{\#g},\mathbb{U},\mathfrak{s}_0)\cong\mathbb{Z}_2[V_{\mathbf{z}}].$$

That the generators for the two maximal subsets agree is obvious, from the diagram we've constructed. □

We now consider the case that all components except for one of the unlink $\mathbb{U}\subseteq(S^1\times S^2)^{\#g}$ have exactly two basepoints, and that the remaining component of $\mathbb{U}$ has exactly four basepoints. These will appear when we define maps associated to bands.

Lemma 3.2. *Suppose that $\mathbb{U}$ is an unlink in $(S^1\times S^2)^{\#g}$ such that all components of $\mathbb{U}$ except for one have exactly two basepoints, and one component of $\mathbb{U}$ has exactly four basepoints. Let z and z' denote the two $\mathbf{z}$-basepoints which are on the same link component. Suppose $\mathfrak{P}_0=\mathbf{w}\cup(\mathbf{z}/(z=z'))$ where $\mathbf{z}/(z=z')$ is the set of $\mathbf{z}$-basepoints, modulo the relation that $z=z'$, and σ_0 is the natural map from $\mathbf{w}\cup\mathbf{z}$ to $\mathfrak{P}_0$. Then we have an isomorphism*

$$\mathcal{H}FL^-_{\max,\mathbf{w}}((S^1\times S^2)^{\#g},\mathbb{U},\sigma_0,\mathfrak{P}_0,\mathfrak{s}_0)\cong\mathbb{Z}_2[V_{\mathbf{z}}]/(V_z-V_{z'}).$$

Similarly if w and w' are the two $\mathbf{w}$ basepoints on this link component, and $(\sigma_0',\mathfrak{P}_0')$ is the coloring with $\mathfrak{P}_0'=(\mathbf{w}/(w=w'))\cup\mathbf{z}$ and σ_0' is the obvious map, we have

$$\mathcal{H}FL^-_{\max,\mathbf{z}}((S^1\times S^2)^{\#g},\mathbb{U},\sigma_0',\mathfrak{P}_0',\mathfrak{s}_0)\cong\mathbb{Z}_2[U_{\mathbf{w}}]/(U_w-U_{w'}).$$

Proof. Again we will only prove the claim about $\mathcal{H}FL^-_{\max,\mathbf{w}}$, since the other one follows identically, by simply switching the w's and z's. This is a modification of the proof of Lemma 3.1. Let $\mathbb{U}_0$ denote the unlink which has all of the components of $\mathbb{U}$ which have exactly two basepoints. Let U denote the component of $\mathbb{U}$ which has four basepoints. Pick a diagram $(\Sigma_g,\boldsymbol{\alpha}_0,\boldsymbol{\beta}_0,\mathbf{w}_0,\mathbf{z}_0)$ for $\mathbb{U}_0$ as in the proof of Lemma 3.1 (in the case that $\mathbb{U}_0=\varnothing$, simply pick a diagram for $(S^1\times S^2)^{\#g}$ which has no basepoints). Now pick a point in the region of

large area, and insert the diagram shown in Figure 3.1 (in the case that $\mathbb{U}_0 = \varnothing$, we insert the diagram as in Figure 3.1, but without the outer α- and β-curves). Let $\mathcal{H} = (\Sigma_g, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ denote the resulting diagram.

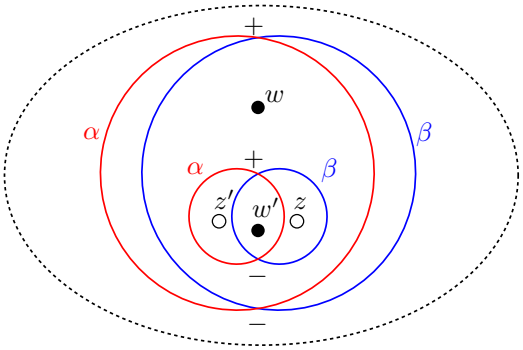

FIGURE 3.1. In the case that $\mathbb{U}_0 \neq \varnothing$, we insert this subdiagram into a diagram $(\Sigma_g, \boldsymbol{\alpha}_0, \boldsymbol{\beta}_0, \mathbf{w}_0, \mathbf{z}_0)$ for $((S^1 \times S^2)^{\#g}, \mathbb{U}_0)$, to get a diagram for $\mathbb{U} = \mathbb{U}_0 \cup U$. The component U has basepoints z, w, z' and w'. If $\mathbb{U}_0 = \varnothing$, we insert the above diagram, but with the outer α- and β-curves omitted.

The same argument that we used in the proof of Lemma 3.1, shows that sums of elements of the form $\Theta \cdot V_{\mathbf{z}}^J$ generate the subset of $\mathcal{C}FL^-(\mathcal{H}, \mathfrak{s})$ of top $\mathrm{gr}_{\mathbf{w}}$-homological degree. Also as before, we note that none of these elements are boundaries.

However, it is important to note that on the uncolored module $\mathcal{C}FL^-$, the element Θ is not a cycle. Instead $\partial\Theta = (V_z + V_{z'})\tau$, for some element τ, as is easily seen in Figure 3.1. On the other hand, the coloring $(\sigma_0, \mathfrak{P}_0)$ identifies V_z and $V_{z'}$, so Θ becomes a cycle in the colored chain complex. Hence

$$\mathcal{H}FL^-_{\max,\mathbf{w}}(\mathcal{H}, \sigma_0, \mathfrak{P}_0, \mathfrak{s})$$

is exactly

$$\mathcal{C}FL^-_{\max,\mathbf{w}}(\mathcal{H}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}) = \mathcal{C}FL^-_{\max,\mathbf{w}}(\mathcal{H}, \mathfrak{s})/(V_z - V_{z'}),$$

which is generated by elements of the form $V_{\mathbf{z}}^J \cdot \Theta$, and has the single relation that $V_z - V_{z'} = 0$. □

Remark 3.3. We will often make use of the existence of an element Θ distinguished by the property of generating a maximally graded subset. For example if $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple and $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ represents an unlink in $(S^1 \times S^2)^{\#g}$, then the map F defined by

$$F(\mathbf{x}) = F_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma},\mathfrak{s}}(\Theta \otimes \mathbf{x})$$

where $F_{\boldsymbol{\alpha}\boldsymbol{\beta}\boldsymbol{\gamma},\mathfrak{s}}$ counts holomorphic triangles, is a well defined, $\mathfrak{P}$-filtered map up to a choice of Θ. It's easy to see that if $\eta \in \mathcal{C}FL^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, then adding $\partial\eta$ to Θ induces an equivariant, filtered chain homotopy, and hence the map F is well defined up to equivariant, filtered chain homotopy as long as the element Θ is defined on homology in $\mathcal{H}FL^-$. We will use such triangle maps to define maps corresponding to surgery on a link, as well as adding bands between components of a link.

3.2. **Quasi-stabilization.** In [Zem16], the author described maps associated to adding or removing adjacent basepoints, w and z along a link component. The operation had already been described in [MO10], but a proof that the operation yielded natural chain maps hadn't been completed. In [Zem16], we colored the the curved complexes $\mathcal{C}FL^\circ$ to get actual chain complexes, which we denoted by CFL°_{UV}.

If w follows z, and $\sigma : \mathbf{w} \cup \mathbf{z} \to \mathfrak{P}$ is a coloring which is extended by $\sigma' : \mathbf{w} \cup \mathbf{z} \cup \{w, z\} \to \mathfrak{P}$, and such that z is given the same color as the other $\mathbf{z}$-basepoint adjacent to w, then there are filtered chain maps

$$S^+_{w,z} : \mathcal{C}FL^\circ(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S^-_{w,z} : \mathcal{C}FL^\circ(Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}, \sigma', \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, L, \mathbf{w}, \mathbf{z}, \sigma, \mathfrak{P}, \mathfrak{s}),$$

which are well defined on the coherent chain homotopy type invariant. If z comes after w, there are maps $S^+_{z,w}$ and $S^-_{z,w}$ defined analogously. Note that the conditions on the coloring were automatic in [Zem16], since in that setting all $V_{\mathbf{z}}$ variables on each link component were identified.

The above statement follows from [Zem16, Theorem A], though in [Zem16], the result was stated for the version of link Floer homology CFL°_{UV}.

We will shortly describe the necessary modifications to the proof in [Zem16]. Additionally, we will describe maps

$$T^+_{w,z} : \mathcal{C}FL^\circ(Y,L,\mathbf{w},\mathbf{z},\sigma,\mathfrak{P},\mathfrak{s}) \to \mathcal{C}FL^\circ(Y,L,\mathbf{w}\cup\{w\},\mathbf{z}\cup\{z\},\sigma',\mathfrak{P},\mathfrak{s})$$

and

$$T^-_{w,z} : \mathcal{C}FL^\circ(Y,L,\mathbf{w}\cup\{w\},\mathbf{z}\cup\{z\},\sigma',\mathfrak{P},\mathfrak{s}) \to \mathcal{C}FL^\circ(Y,L,\mathbf{w},\mathbf{z},\sigma,\mathfrak{P},\mathfrak{s}),$$

defined when w has the same color as the other $\mathbf{w}$ basepoint adjacent to z.

We first outline the quasi-stabilization operation.

Suppose we are given a Heegaard diagram $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathbf{z})$ for $\mathbb{L} = (L,\mathbf{w},\mathbf{z})$. Given a choice of new, adjacent basepoints w and z, such that w comes after z, we let A denote the component of $\Sigma\setminus\boldsymbol{\alpha}$ containing the basepoints adjacent to w and z. We pick a point $p \in A\setminus(\boldsymbol{\alpha}\cup\boldsymbol{\beta}\cup\mathbf{w}\cup\mathbf{z})$ and a choice of new α_s curve in A passing through p. We then form the "special connected sum" of $\mathcal{H}$ with a diagram $(S^2,\alpha_0,\beta_0,w,z)$ at a point $p_0\in\alpha_0$ chosen, as in Figure 3.2.

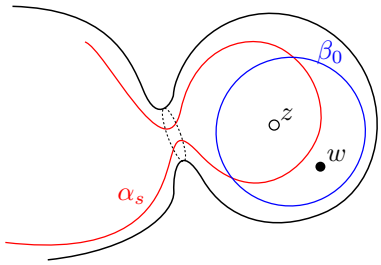

FIGURE 3.2. A quasi-stabilized diagram. We stretch the almost complex structure along the dashed curve.

Note that there are two new intersection points, which are distinguished by the relative Maslov grading. Of course we now have two Maslov gradings, namely $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$. Let $\theta^+_{\mathbf{w}}$ denote the higher generator with respect to $\mathrm{gr}_{\mathbf{w}}$ and let $\theta^+_{\mathbf{z}}$ denote the higher generator with respect to $\mathrm{gr}_{\mathbf{z}}$. By inspection, we have that

$$\theta^\pm_{\mathbf{w}} = \theta^\mp_{\mathbf{z}}.$$

In [Zem16, Prop. 5.3], the quasi-stabilized differential is computed:

Proposition 3.4. *If $\mathcal{H}$ is a diagram for $\mathbb{L} = (L,\mathbf{w},\mathbf{z})$, and $\overline{\mathcal{H}}$ denotes the quasi-stabilized diagram, then for appropriately stretched almost complex structures, there is an identification of differentials*

$$\partial_{\overline{\mathcal{H}}} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_{w'} \\ V_z + V_{z'} & \partial_{\mathcal{H}} \end{pmatrix},$$

where w' and z' are basepoints on $\mathbb{L}$ adjacent to w and z. Here the matrix is in terms of the generators $\theta^\pm_{\mathbf{w}}$ (i.e. $\theta^+_{\mathbf{w}}$ is corresponds to the first row and column and $\theta^-_{\mathbf{w}}$ corresponds to the second row and column).

In particular, if $(\sigma,\mathfrak{P})$ is a coloring of the basepoints of $\mathbb{L}$ which sends z and z' to the same color, then we can define maps by the formulas

$$S^+_{w,z}(\mathbf{x}) = \mathbf{x}\otimes\theta^+_{\mathbf{w}}$$

and

$$S^-_{w,z}(\mathbf{x}\otimes\theta^+_{\mathbf{w}}) = 0, \qquad \text{and} \qquad S^-_{w,z}(\mathbf{x}\otimes\theta^-_{\mathbf{w}}) = \mathbf{x},$$

and these maps are chain maps.

We similarly, if $(\sigma, \mathfrak{P})$ is a coloring which gives w and w' the same color, then we can define chain maps

$$T^+_{w,z}(\mathbf{x}) = \mathbf{x} \otimes \theta^+_{\mathbf{z}}$$

and

$$T^-_{w,z}(\mathbf{x} \otimes \theta^+_{\mathbf{z}}) = 0, \qquad \text{and} \qquad T^-_{w,z}(\mathbf{x} \otimes \theta^-_{\mathbf{z}}) = \mathbf{x}.$$

Under these coloring assumptions, the proof of [Zem16, Theorem A] carries over without modification to show that the maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ each commute up to filtered, equivariant chain homotopy with the change of diagrams maps $\Phi_{\mathcal{H}\to\mathcal{H}'}$. The key computation is a challenging triangle map computation, which is stated in this paper as Theorem 5.3.

In the setting of [Zem16], the maps $T^{\pm}_{w,z}$ correspond to $\Psi_z S^+_{w,z}$ and $S^-_{w,z}\Psi_z$. In our setting, we have the following:

Lemma 3.5. *Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link and w and z are new basepoints, in a component of $L \setminus (\mathbf{w} \cup \mathbf{z})$. Say that w and z are between w' and z', two basepoint of $\mathbb{L}$. If w and w' are given the same color, and z and z' are given the same color, then*

$$T^+_{w,z} \simeq \Psi_z S^+_{w,z}, \qquad \text{and} \qquad T^-_{w,z} \simeq S^-_{w,z}\Psi_z.$$

Similarly

$$S^+_{w,z} \simeq \Phi_w T^+_{w,z}, \qquad \text{and} \qquad S^-_{w,z} \simeq T^-_{w,z}\Phi_w.$$

Proof. This is an easy computation following from the definitions of the maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$, as well as Φ_w and Ψ_z, which can be easily computed from Proposition 3.4. Note that the requirement on the coloring guarantees that all of the above maps are chain maps. □

3.3. **The maps Φ_w and Ψ_z and relations amongst them.** Two additional maps which appear in the link cobordism maps are the maps Φ_w and Ψ_z. These correspond to "formally differentiating the differential with respect to the variables U_w or V_z". They are defined by the formulas

$$\Phi_w(\mathbf{x}) = U_w^{-1} \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} n_w(\phi)\#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y},$$

and

$$\Psi_z(\mathbf{x}) = V_z^{-1} \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} n_z(\phi)\#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y}.$$

According to [Zem16, Lemma 3.2], these maps commute with the change of diagrams maps, up to equivariant, filtered chain homotopy. We now summarize several important properties of the maps Φ_w and Ψ_z in the following lemmas.

By differentiating the expression in Lemma 2.9 once, with respect to a single variable, we have the following:

Lemma 3.6 ([Zem16, Lemma 3.1])**.** *If z is a basepoint, which is adjacent to basepoints w and w', then we have*

$$\partial\Psi_z + \Psi_z\partial = U_w + U_{w'}.$$

If w is a basepoint which is adjacent to z and z', then we have that

$$\partial\Phi_w + \Phi_w\partial = V_z + V_{z'}.$$

Note that Φ_w and Ψ_z are not chain maps for the trivial coloring. In fact, they induce chain maps iff the variables U_w and $U_{w'}$ are identified, or the variables V_z and $V_{z'}$ are identified, respectively. Notice that this corresponds exactly to the condition that the coloring of the basepoints agree with a coloring of the appropriate decorated surface in Figure 1.2.

By differentiating the expression in Lemma 2.9 twice, the commutator of various Φ_w and Ψ_z is computed in [Zem16, Lemmas 9.1 and 9.2]. Those two lemmas can be summarized as follows:

Lemma 3.7. *If w, w', z and z' are basepoints on $\mathbb{L}$, we have*

$$\Phi_w\Phi_{w'} + \Phi_{w'}\Phi_w \simeq 0,$$

and

$$\Psi_z\Psi_{z'} + \Psi_{z'}\Psi_z \simeq 0.$$

We also have

$$\Phi_w\Psi_z + \Psi_z\Phi_w \simeq N(w,z),$$

where $N(w,z)$ is the number of components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which have both w and z as boundary.

We make the following remark (compare [Zem15b, Corollary 14.11]):

Remark 3.8. If $\mathcal{H}$ is a diagram for $(Y, \mathbb{L})$, then on the uncolored curved complex $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$, the maps Φ_w are nonequivariantly chain homotopic to zero. The chain homotopy is given by

$$\Phi_w = \frac{d}{dU_w} \circ \partial + \partial \circ \frac{d}{dU_w}.$$

The chain homotopy may not persist if we color and set some of the $U_{\mathbf{w}}$-variables equal to each other. Notice also that we work in the category of filtered, equivariant homomorphisms up to *filtered, equivariant* chain homotopies, and hence Φ_w is not in general the zero morphism, despite the above chain homotopy.

Similarly on the uncolored complex, we have

$$\Psi_z = \frac{d}{dV_z} \circ \partial + \partial \circ \frac{d}{dV_z}$$

on $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$. Again this chain homotopy may not persist after we color the complexes.

The maps Φ_w and Ψ_z are homotopy differentials (cf. [Zem16, Lemma 9.7]):

Lemma 3.9. *The endomorphisms Φ_w and Ψ_z satisfy*

$$\Phi_w^2 \simeq 0, \qquad \text{and} \qquad \Psi_z^2 \simeq 0,$$

through filtered, equivariant chain homotopies.

Proof. This is a purely algebraic. We will show that $\Phi_w^2 \simeq 0$; the proof for Ψ_z will be identical. Write $\partial = \sum_{k=0}^{\infty} \partial_k U_w^k$, where ∂_i does not involve U_w. Using this notation, we have

$$\Phi_w = \sum_{k=0}^{\infty} k\partial_k U^{k-1}, \qquad \text{and} \qquad \partial^2 = \sum_{k\geq 0} \sum_{n+m=k} \partial_n\partial_m U^{n+m}.$$

Applying Lemma 2.9 for any $k \geq 2$, we have that

$$\sum_{m+n=k} \partial_n\partial_m = 0. \tag{1}$$

Define

$$H = \sum_{k=0}^{\infty} \frac{k(k-1)}{2} \partial_k U_w^{k-2}.$$

One now simply computes that

$$\Phi_w^2 + (\partial H + H\partial) = \sum_{m,n\geq 0} \frac{(m+n)(m+n-1)}{2} \partial_n\partial_m U_w^{n+m-2} = \sum_{k=0}^{\infty} \frac{k(k-1)}{2} U_w^{k-2} \sum_{n+m=k} \partial_n\partial_m,$$

which we see vanishes, because the $k = 0$ and $k = 1$ terms vanish trivially and the terms for $k \geq 2$ vanish by Equation (1). □

Remark 3.10. The fact that $\Phi_w^2 \simeq 0$ and $\Psi_z^2 \simeq 0$ is a purely algebraic fact arising from our computation of ∂^2 in Lemma 2.9, and presents no real obstruction to a complex coming from Heegaard Floer homology. Suppose $\mathcal{R}$ is a polynomial ring (over $\mathbb{Z}$ or $\mathbb{Z}_2$) and C is a free module over $\mathcal{R}$ with a chosen free basis. Suppose E is an $\mathcal{R}$-equivariant endomorphism and U is a variable of the polynomial ring. If $E^2 = F_0 + F_1 U$, for endomorphisms F_0 and F_1 with no U terms (e.g. if $E^2 = 0$), then

$$(E')^2 + EH + HE = 0,$$

where E' is the formal derivative of E with respect to U, and H is a map defined like the chain homotopy in the previous lemma.

3.4. **Relations between the quasi-stabilization maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ and the maps Φ_w and Ψ_z.** In this section, we summarize some of the important relations satisfied by the quasi-stabilization maps and the maps Φ_w and Ψ_z. Similar statements can be found in [Zem16].

By considering the quasi-stabilized differential in Proposition 3.4, one can derive the following:

Lemma 3.11 ([Zem16, Lemma 9.3])**.** *If w is not the only $\mathbf{w}$-basepoint on the link component containing w (so that $S^-_{w,z,\mathbf{w}}$ can be defined), we have that, as endomorphisms of $\mathcal{C}FL(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$ for a fixed coloring $(\sigma,\mathfrak{P})$,*

$$\Phi_w \simeq S^+_{w,z}S^-_{w,z}, \qquad \text{or} \qquad \Phi_w \simeq S^+_{z,w}S^-_{z,w}$$

for z an adjacent $\mathbf{z}$-basepoint.

Proof. This can be proven by explicit examination of the quasi-stabilized differential in Proposition 3.4, as well as the formulas for $S^{\pm}_{w,z}$. □

Note that Φ_w is defined as an endomorphism of a particular complex. Nonetheless the maps $S^+_{w,z}S^-_{w,z}$ can also be defined as maps between complexes when w is given a different coloring on the first complex than the second. In the latter case, a map Φ_w can no longer be defined using the formula we used above. See also [Zem15b, Theorem 14.1].

Symmetrically, we have the following:

Lemma 3.12. *Suppose z is not the only $\mathbf{z}$-basepoint on the link component containing z (so that $T^-_{w,z}$ can be defined), then as endomorphisms of $\mathcal{C}FL(Y,\mathbb{L},\sigma,\mathfrak{P},\mathfrak{s})$, for a fixed coloring $(\sigma,\mathfrak{P})$, we have that*

$$\Psi_z = T^+_{w,z}T^-_{w,z}$$

if w comes immediately after z. If w comes immediately before z, we have

$$\Psi_z = T^+_{z,w}T^-_{z,w}.$$

Proof. The proof is identical to the proof of Lemma 3.11. □

Additionally we have the following:

Lemma 3.13. *The quasi-stabilization maps satisfy*

$$S^-_{w,z}S^+_{w,z} \simeq 0 \qquad \text{and} \qquad T^-_{w,z}T^+_{w,z} \simeq 0,$$

as well as

$$T^-_{w,z}S^+_{w,z} \simeq \mathrm{id}, \qquad \text{and} \qquad S^-_{w,z}T^+_{w,z} \simeq \mathrm{id}\,.$$

Proof. These are all just easy applications of the formulas involved. □

Remark 3.14. The latter two formulas should be thought of as an analog to [Zem16, Lemma 9.6]. In that context, we had $S^-_{w,z}\Psi_z S^+_{w,z} \simeq \mathrm{id}$. In our context, $T^-_{w,z} \simeq S^-_{w,z}\Psi_z$. Note that if add basepoints (w,z) in between w' and z', and w' is colored the same as w, and z' is colored the same as z, then using Lemma 3.5, we see that indeed we still have $S^-_{w,z}\Psi_z S^+_{w,z} \simeq \mathrm{id}$.

In [Zem16], we show the following result about quasi-stabilization maps commuting:

Theorem 3.15 ([Zem16, Theorem 8.1])**.** *The quasi-stabilization maps $S^{\pm}_{w,z}$ commute. That is if (w,z) and (w',z') are disjoint pairs of basepoints, then*

$$S^{\circ}_{w,z}S^{\circ'}_{w',z'} \simeq S^{\circ'}_{w',z'}S^{\circ}_{w,z},$$

for any choice of $\circ,\circ' \in \{+,-\}$. Commutation still holds if both of the pairs of basepoints (w,z) and (w',z') appear with the opposite ordering (i.e. if z appears after w, and we use the maps $S^{+}_{z,w}$, instead).

The proof can be found in [Zem16, Section 8]. The proof in [Zem16] of the previous theorem also carries over without change if we replace both **w**-quasi-stabilizations by **z**-quasi-stabilizations:

Lemma 3.16. *The quasi-stabilization maps $T^{\pm}_{w,z}$ commute. That is if (w,z) and (w',z') are disjoint pairs of basepoints, then*

$$T^{\circ}_{w,z}T^{\circ'}_{w',z'} \simeq T^{\circ'}_{w',z'}T^{\circ}_{w,z},$$

for any choice of $\circ,\circ' \in \{+,-\}$. Commutation still holds if either or both of the pairs of basepoints (w,z) and (w',z') appears with the opposite ordering (i.e. if z appears after w, and we use the maps $T^{+}_{z,w}$, instead).

The previous two results don't immediately tell us whether $S^{\pm}_{w,z}$ and $T^{\pm}_{w',z'}$ commute. In light of Lemma 3.13, they will not always commute. This will be important to our results, so we now consider it.

Proposition 3.17. *Suppose (w,z) are (w',z') are disjoint pairs of adjacent basepoints, with w following z and w' following z'. If w' is not adjacent to z, then*

$$T^{\circ}_{w,z}S^{\circ'}_{w',z'} \simeq S^{\circ'}_{w',z'}T^{\circ}_{w,z},$$

for any choice of $\circ,\circ' \in \{+,-\}$. The same result holds if z follows w and z' follows w', with $T^{\circ}_{w,z}$ changed to $T^{\circ}_{z,w}$ and $S^{\circ'}_{w',z'}$ changed to $S^{\circ'}_{z',w'}$, as long as z and w' are not adjacent.

Proof. We will consider only the the first statement, since the second follows by a straightforward modification. We refer the reader to [Zem16, Section 8], since our argument here will be a straightforward modification. We take a diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for the unstabilized link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and let $p, p' \in \Sigma \setminus \boldsymbol{\alpha}$ be the two points at which we perform the quasi-stabilization. We form the quasi-stabilized diagram $\overline{\mathcal{H}}$, as in Figure 8.1. Note that the formulas for the two quasi-stabilization maps obviously commute, but it's not clear whether we can use the same almost complex structure to compute both.

There are two cases to consider. The first case is that p and p' are in separate components of $\Sigma \setminus \boldsymbol{\alpha}$, and the second is that p and p' occur in the same component of $\Sigma \setminus \boldsymbol{\alpha}$. The case that p and p' are in separate components of $\Sigma \setminus \boldsymbol{\alpha}$ corresponds to (w,z) and (w',z') being in separate components of $L \setminus (\mathbf{w} \cup \mathbf{z})$. There are two necks to stretch, and we let $\mathcal{J}$ denote gluing data to stretch on both necks. According to [Zem16, Lemma 8.2], for all sufficiently large T_1, T_1', T_2, and T_2', the change of almost complex structure map

$$\Phi_{\mathcal{J}(T_1,T_1') \to \mathcal{J}(T_2,T_2')} = \mathrm{id},$$

under the obvious identification of the two complexes. Hence in this case, the maps $T^{\pm}_{w,z}S^{\pm}_{w',z'} \simeq S^{\pm}_{w',z'}T^{\pm}_{w,z}$ obviously commute.

We now consider the second case, where p and p' are in the same component of $\Sigma \setminus \boldsymbol{\alpha}$. This case corresponds to (w,z) and (w',z') being adjacent on $L \setminus (\mathbf{w} \cup \mathbf{z})$. Since w' is not adjacent to z, we note that this implies that the basepoints appear in the following order z, w, z' then w'. The change of almost complex structure maps $\Phi_{\mathcal{J}(T_1,T_1') \to \mathcal{J}(T_2,T_2')}$ are computed in Lemma 8.3. By the proof of that Lemma, if $F = \Phi_{\mathcal{J}(T_1,T_1') \to \mathcal{J}(T_2,T_2')}$ then

$$\begin{aligned}
F(\mathbf{x} \times \theta^{-}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{+}) &= \mathbf{x} \times \theta^{-}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{+} \\
F(\mathbf{x} \times \theta^{-}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{-}) &= \mathbf{x} \times \theta^{-}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{-} + C \cdot \mathbf{x} \times \theta^{+}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{+} \\
F(\mathbf{x} \times \theta^{+}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{+}) &= \mathbf{x} \times \theta^{+}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{+} \\
F(\mathbf{x} \times \theta^{+}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{-}) &= \mathbf{x} \times \theta^{+}_{\mathbf{z}} \times (\theta'_{\mathbf{w}})^{-},
\end{aligned}$$

for some C (which is not independent of T_i and T_i').

It is now easy (though somewhat tedious) algebra to check that the maps commute. □

Remark 3.18. The reader is encouraged to check that the strategy used in the previous proof fails to show that $S^-_{w,z}T^+_{w',z'}$ and $T^+_{w',z'}S^-_{w,z}$ are chain homotopic if z' and w are adjacent and we have the following ordering of basepoints: z, w, z' then w'. One should compare this to the interpretation of the maps in terms of the dividing sets shown in Figure 1.2.

We now give simple repackaging of [Zem16, Lemma 9.4] and Lemma [Zem16, 9.5].

Suppose that A is a subarc of L, between two $\mathbf{w}$-basepoints. We do not exclude the case that A consists of a single point (which we treat as the empty arc), or an entire link component of L. We define

$$\Psi_A = \sum_{z\in A} \Psi_z.$$

Note that Ψ_z is a chain map iff the endpoints of A share the same color.

Lemma 3.19. *If w and z are new, adjacent basepoints, we have*

$$\Psi_A S^{\pm}_{w,z} + S^{\pm}_{w,z}\Psi_A \simeq 0.$$

Similarly

$$\Psi_A T^{\pm}_{w,z} + T^{\pm}_{w,z}\Psi_A \simeq 0.$$

Proof. Let $\mathfrak{P} = \mathbf{w}\cup\mathbf{z}\cup\{w\cup z\}$, and let $\sigma : (\mathbf{w}\cup\mathbf{z}) \to \mathfrak{P}$ denote the natural inclusion, and let $\sigma' : \mathbf{w}\cup\mathbf{z}\cup\{w\cup z\} \to \mathfrak{P}$ denote the identity. Let $\mathcal{H}^-$ denote a diagram for $(L, \mathbf{w}, \mathbf{z})$, and let $\mathcal{H}^+$ denote it's quasi-stabilization resulting from adding the basepoints w and z. Consider the colored complexes $\mathcal{C}FL^\circ(\mathcal{H}^-, \sigma, \mathfrak{P}, \mathfrak{s})$ and $\mathcal{C}FL^\circ(\mathcal{H}^+, \sigma', \mathfrak{P}, \mathfrak{s})$. We note that using Proposition 3.4, we can write

$$\partial_{\mathcal{H}^\pm} S^{\pm}_{w,z} + S^{\pm}_{w,z}\partial_{\mathcal{H}^\mp} + (V_z + V_{z'})F = 0, \tag{2}$$

where F is a map which has does not involve the variables in $U_{\mathfrak{P}}$. Here z' is the $\mathbf{z}$-basepoint adjacent to z. We view these as matrices with coefficients in $\mathbb{Z}_2[U_{\mathfrak{P}}]$. Let

$$D_A = \sum_{z\in A} \frac{d}{dV_z}.$$

Noting that $\frac{d}{dV_z}\partial_{\mathcal{H}^-} = 0$, we have that $D_A(\partial_{\mathcal{H}^\circ}) = \Psi_A$ for both $\circ = +$ and $\circ = -$. We also note that on the quasi-stabilized link, either z and z' are in A, or neither are in A, as w is not an endpoint of A. Hence

$$D_A((V_z + V_{z'})F) = 0,$$

as is easily computed using the Leibniz rule, separately, in the case that $z, z' \in A$ and the case that $z, z' \notin A$, using also the fact that $D_A F = 0$, as the matrix expression for F involves none of the variables in A, upon direct inspection. Thus applying D_A to Equation (2) yields

$$\Psi_A S^{\pm}_{w,z} + S^{\pm}_{w,z}\Psi_A \simeq 0.$$

The statement involving $T^{\pm}_{w,z}$ follows from similar considerations. □

We can prove the analogous statement for the maps Φ_w. If A is an arc on $(L, \mathbf{w}, \mathbf{z})$ between two $\mathbf{z}$-basepoints, define

$$\Phi_A = \sum_{w\in A} \Phi_w.$$

The map Φ_A is a chain map iff the $\mathbf{z}$-basepoints on the ends share the same color. Analogously to the previous lemma, we have the following:

Lemma 3.20. *If A is an arc on L between two $\mathbf{z}$-basepoints of $(L, \mathbf{w}, \mathbf{z})$, and w and z are two new, adjacent basepoints on L, then*

$$S^{\pm}_{w,z}\Phi_A + \Phi_A S^{\pm}_{w,z} \simeq 0,$$

and

$$T^{\pm}_{w,z}\Phi_A + \Phi_A T^{\pm}_{w,z} \simeq 0.$$

The proof is identical to the proof of Lemma 3.19.

We now highlight several useful special cases of the previous lemmas:

Lemma 3.21 ([Zem16, Lemma 9.4]). *If w and w' are distinct basepoints, w follows z, and the $\mathbf{z}$-basepoints adjacent to w' have the same color, then*

$$S^{\pm}_{w,z}\Phi_{w'} + \Phi_{w'}S^{\pm}_{w,z} \simeq 0.$$

A similar statement holds for $S^{\pm}_{z,w}$ if z comes after w.

If z' is not adjacent to a basepoint w, and z comes immediately before w, then

$$S^{\pm}_{w,z}\Psi_{z'} + \Psi_{z'}S^{\pm}_{w,z} \simeq 0.$$

Proof. The first relation is obtained by picking an arc A between the two $\mathbf{z}$-basepoints adjacent to w', before quasi-stabilizing. Applying Lemma 3.19 yields the result immediately in the case that A does not contain the basepoints w and z. In the case that A contains the basepoints w and z, we note that the result follows from the additional observation that $S^{-}_{w,z}\Phi_w \simeq 0$ and $\Phi_w S^{+}_{w,z} \simeq 0$.

The second claim follows from the same strategy. □

Lemma 3.22 ([Zem16, Lemma 9.5]). *If z and z' are the basepoints immediately adjacent to w, and w comes after z (and $z \neq z'$), then we have*

$$S^{+}_{w,z}\Psi_{z'} = (\Psi_z + \Psi_{z'})S^{+}_{w,z}$$

and

$$S^{-}_{w,z}(\Psi_z + \Psi_{z'}) = \Psi_{z'}S^{-}_{w,z}.$$

A similar statement holds of the maps $S_{z,w}$, in the case that z comes after w.

Proof. This follows immediately from Lemma 3.19, picking A to be the arc on between the $\mathbf{w}$-basepoints on the unstabilized link adjacent to z'. □

3.5. **Quasi-stabilization and basepoint moving maps I.** In this section we consider the relation between the quasi-stabilization maps, $S^{\pm}_{w,z}, S^{\pm}_{z,w}, T^{\pm}_{w,z}$ and $T^{\pm}_{z,w}$, and certain basepoint moving maps. We will later define four classes of simple dividing sets, which we will denote by I^+, I^-, II^+ and II^-, as shown in Figure 11.1. For a given dividing set as in Figure 11.1, there are two maximal identifications of a subset of the incoming boundary with a subset of the outgoing boundary, respecting the divides. We now prove several results which will imply that cobordism maps we will define in Subsection 11.2 do not depend on the choice of these two identifications.

3.5.1. *Relations for the maps $S^{\pm}_{w,z}$.* Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link, and $z' \in \mathbf{z}$.

Now suppose that w and z are two adjacent basepoints on L which are not in $\mathbf{w} \cup \mathbf{z}$, such that, as one traverses the link according to its orientation, these basepoints appear in the order z', w then z. Let ρ be the diffeomorphism

$$\rho_* : (Y, L, \mathbf{w}, \mathbf{z}) \to (Y, L, \mathbf{w}, (\mathbf{z} \setminus \{z\}) \cup \{z'\})$$

such that $\rho(z') = z$, and is fixed outside of a neighborhood of the arc between z' and z (note that $w, z \notin \mathbf{w} \cup \mathbf{z}$).

Lemma 3.23. *We have a chain homotopy equivalence*

$$S^{+}_{w,z} \simeq S^{+}_{z',w}\rho_*.$$

Proof. Consider the diagram shown in Figure 3.3.

If we stretch along the circle c', we get a complex structure which can be used to compute $S^{+}_{w,z}$. If we change the almost complex structure so that it is stretched along c'', we get a diagram which we can use to $S^{+}_{z',w}$. On the diagram occurring before quasi-stabilization (with only the z' basepoint), the map ρ_* has the effect of just moving z' to z along the surface. Thus it is sufficient to compute the effect of changing the almost complex structure by switching from a β-quasi-stabilization to an α-quasi-stabilization.

Let $\mathcal{H}_0 = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \setminus\{z'\})$ denote the partial Heegaard diagram. The partial diagram $\mathcal{H}_0$ has a distinguished point p, the connected sum point. Pick a generic almost complex structure J_s on $\mathcal{H}_0$, i.e., an almost complex structure on $\Sigma \times [0,1] \times \mathbb{R}$, which is split in a neighborhood of p. In [Zem15b, Definition 6.1] we define "gluing data" $\mathcal{J}$ to be a collection of data, including a choice of almost complex structures J_s on $\Sigma \times [0,1] \times \mathbb{R}$ and J'_s on $S^2 \times [0,1] \times \mathbb{R}$, as well as some auxiliary data which allow one to construct an almost

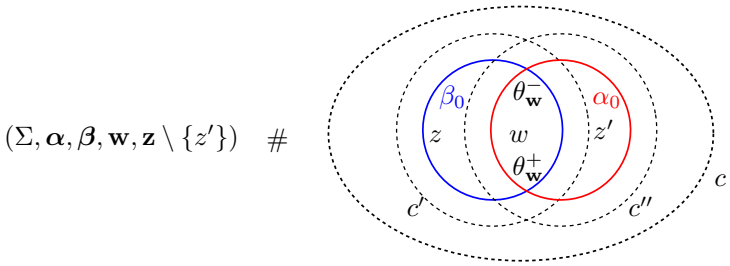

FIGURE 3.3. The diagram we consider in Lemma 3.23. We must consider the change of almost complex structure map as we change from an α-quasi-stabilization (stretched along c'') and a β-quasi-stabilization (stretched along c').

complex structure $\mathcal{J}(T)$, for a choice of necklength T (see [Zem15b, Subsection 6.1] for more details). As $T \to \infty$, the almost complex manifolds

$$((\Sigma \# S^2) \times [0,1] \times \mathbb{R}, \mathcal{J}(T))$$

approach

$$((\Sigma \vee S^2) \times [0,1] \times \mathbb{R}, J_s \vee J'_s).$$

Performing the obvious adaptation, we can consider a collection of gluing data $\mathcal{J}_{c'}$ for stretching along c and c'. We let $\mathcal{J}_{c,c'}(T,T')$ denote the resulting almost complex structure, with neck lengths T and T' inserted along c and c'' respectively. We define a collection of gluing data $\mathcal{J}_{c,c''}$ for stretching along c and c'' analogously.

We make the following claim: *if T is sufficiently large, then for any $T', T'' > 0$ we can pick a dynamic almost complex structure $\widehat{J}$ such that the change of almost complex structure map, computed with the dynamic almost complex structure $\widehat{J}$, takes the form*

$$\Phi_{\mathcal{J}_{c,c'}(T,T') \to \mathcal{J}_{c,c''}(T,T'')} = \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}. \tag{3}$$

Here the first component corresponds to the $\theta^+_{\mathbf{w}}$ generator and the second corresponds to the $\theta^-_{\mathbf{w}}$ generator. We claim the choice of T can be made independent of T' and T''.

Note that, assuming this, the lemma statement follows immediately, since with respect to $\mathcal{J}_{c,c'}(T,T')$ (for sufficiently large T') the map $S^+_{w,z}$ takes the form

$$\mathbf{x} \mapsto \mathbf{x} \times \theta^+_{\mathbf{w}},$$

and with respect to $\mathcal{J}_{c,c''}(T,T'')$ the map $S^+_{z',w}$ takes the same form. Assuming the change of almost complex structure map $\Phi_{\mathcal{J}_{c,c'}(T,T') \to \mathcal{J}_{c,c''}(T,T'')}$ takes the form claimed in Equation (3), we see that $S^+_{w,z}$ and $S^+_{z',w}\rho_*$ are equal.

Hence it is sufficient to show that $\Phi_{\mathcal{J}_{c,c'}(T,T') \to \mathcal{J}_{c,c''}(T,T'')}$ takes the form in Equation (3) for sufficiently large T, independent of T' and T''.

We argue by the following weak limits argument. First let us emphasize that J_s is a cylindrical almost complex structure, and hence admits no nonconstant Maslov index zero disks, by transversality. Suppose that there were representatives u_{T_i}, for dynamic almost complex structures $\widehat{\mathcal{J}}_i$, which are cylindrical on a subset $(\Sigma \setminus B) \cup (S^1 \times [-T_i, T_i]) \subseteq \Sigma$, for arbitrarily large choices of T_i. Without loss of generality suppose that u_{T_i} all represent the same homology class ϕ. Write

$$\phi = \phi_\Sigma \# \phi_{S^2} \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times y),$$

where ϕ_Σ is a homology class on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and ϕ_{S_2} is a homology class on (S^2, α_0, β_0). One computes easily that

$$\mu(\phi) = \mu(\phi_1) + \Delta_{\mathbf{w}}(x,y) + 2n_w(\phi_{S^2})$$

where $\Delta_{\mathbf{w}}(x,y)$ denotes the drop in grading from x to y, with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading. By our assumptions on $\widehat{\mathcal{J}}_i$, we can extracting a weak limit to a broken curve representing ϕ_1, and hence we conclude that $\mu(\phi_1) \geq 0$. Hence for $\Delta(x,y) \geq 0$ (i.e. the diagonal and lower diagonal entries) we conclude that for a representative to exist, we must have

$$\mu(\phi) = \Delta_{\mathbf{w}}(x,y) = n_w(\phi_{S^2}) = 0.$$

Since ϕ has a broken representative for J_s, a cylindrical almost complex structure, we conclude by transversality that ϕ must be a constant disk. On the other hand, constant disks are always counted. Since the lower left entry of the transition matrix corresponds to $\Delta_{\mathbf{w}}(x,y) = 1$, no disks are counted by that component, since the existence of a broken representative of ϕ_1 would force

$$\mu(\phi) = \mu(\phi_1) + \Delta(x,y) + 2n_w(\phi_{S^2}) \geq 1.$$

Hence the entries of the transition matrix are as claimed. □

We now consider the relation that will be useful for type I^- cylindrical cobordisms. Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link with a basepoint $z_0 \in \mathbf{z}$. Let z, w, w', z' be new basepoints in the two components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ adjacent to z_0. Assume that these basepoints appear in the order z, w, z_0, w' then z', according to the orientation of the link. Note that there is a well defined diffeomorphism

$$\rho' : (Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}) \to (Y, L, \mathbf{w} \cup \{w'\}, \mathbf{z} \cup \{z'\})$$

which is the identity outside of a neighborhood of the arc from z' to z, and satisfies

$$\rho'(z) = z_0, \qquad \rho'(w) = w', \qquad \text{and} \qquad \rho'(z_0) = z'.$$

Analogously to Lemma 3.23, we have the following:

Lemma 3.24. *We have that*

$$S^-_{w,z} \simeq S^-_{z',w'}\rho'_*.$$

Proof. The proof follows easily from the same computation of the change of almost complex structure maps performed in Lemma 3.23 □

3.5.2. *Relations for the maps $T^\pm_{w,z}$.* We now wish to prove the analogs of Lemmas 3.23 and 3.24, but for the maps $T^\pm_{w,z}$.

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link with a basepoint $w' \in \mathbf{w}$. Suppose that z and w are two basepoints in $L \setminus (\mathbf{w} \cup \mathbf{z})$ which appear in the order w', z then w, with no other basepoints in between these basepoints. There is a well defined diffeomorphism

$$\rho : (Y, L, \mathbf{w}, \mathbf{z}) \to (Y, L, (\mathbf{w} \setminus \{w'\}) \cup \{w\}, \mathbf{z})$$

which is the identity outside of a neighborhood of the arc from w to w', and satisfies

$$\rho(w') = w.$$

We have the following:

Lemma 3.25. *We have that*

$$T^+_{w,z} \simeq T^+_{z,w'}\rho_*.$$

Proof. As with the proof of Lemma 3.23, the key is a change of almost complex structures map. The analogous diagram is shown in Figure 3.4.

Again we just need to compute the matrix form of each map and show they agree. Both $T^+_{w,z}$ and $T^+_{z,w'}$ take the matrix form $\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$, where now we put the generator of higher $\mathrm{gr}_{\mathbf{z}}$ grading as the first generator. Note that if $\phi = \phi_\Sigma \# \phi_{S^2} \in \pi_2(\mathbf{x} \times x, \mathbf{y} \times y)$, is a homology disk then

$$\mu(\phi) = \mu(\phi_\Sigma) + 2n_{z'}(\phi_{S_2}) + \Delta_{\mathbf{z}}(x,y).$$

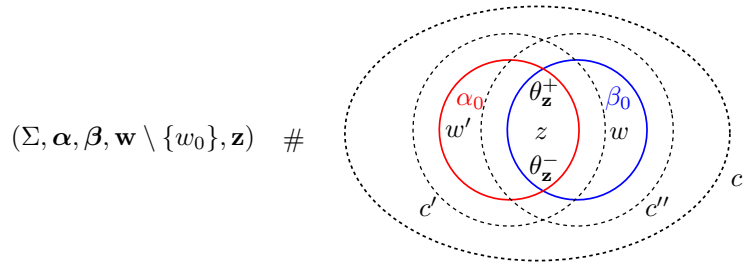

FIGURE 3.4. The diagram we consider in Lemma 3.25.

Thus to show the equality of the lemma statement, we need to compute the change of almost complex structures map and show that it preserves $\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$. The weak limits argument from Lemma 3.23 now applies to show that the relevant change of almost complex structure map is now

$$\begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix},$$

which clearly commutes with $\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$.[3] □

We now consider the analogous statement for model cylindrical cobordisms of type II^-. Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link with a basepoint $w_0 \in \mathbf{w}$. Let w', z', z, w be new basepoints in the two components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ adjacent to z_0. Assume that these basepoints appear in the order w', z', w_0, z then w, according to the orientation of the link. Note that there is a well defined diffeomorphism

$$\rho' : (Y, L, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\}) \to (Y, L, \mathbf{w} \cup \{w'\}, \mathbf{z} \cup \{z'\})$$

which is the identity outside of a neighborhood of the arc from w' to w, and satisfies

$$\rho'(w) = w_0, \qquad \rho'(z) = z', \qquad \text{and} \qquad \rho'(w_0) = w'.$$

Analogously to Lemma 3.24, we have the following:

Lemma 3.26. *We have that*

$$T^-_{w,z} \simeq T^-_{z',w'} \rho'_*.$$

Proof. The proof is a straightforward modification of Lemma 3.25. □

3.6. **Quasi-stabilization and basepoint moving maps II.** In [Zem16], we proved that if (w, z) and (w', z') are adjacent pairs of basepoints, then the map induced by the diffeomorphism ϱ_0 resulting from moving the pair (w, z) to (w', z') is chain homotopic to

$$(\varrho_0)_* \simeq S^-_{w,z} \Psi_{z'} S^+_{w',z'}. \tag{4}$$

A proof can be found in [Zem16, Lemma 10.1]. In this section, we will give a new proof of the above formula, and prove some related results.

Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link and $z \in \mathbf{z}$. Let w and z' be new basepoints, in one of the components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ adjacent to z. There is a well defined diffeomorphism $\phi_{z \to z'}$ obtained from just moving z to z' and fixing all the other basepoints. Suppose that the basepoints appear in the order z, w, then z'.

Lemma 3.27. *The map induced by $\phi_{z \to z'}$ takes the form*

$$(\phi_{z \to z'})_* \simeq T^-_{w,z} S^+_{z',w} \simeq S^-_{w,z} T^+_{z',w}.$$

[3]Note that this the matrix decomposition for the change of almost complex structure map is now in the $\mathrm{gr}_{\mathbf{z}}$ grading, whereas in Lemma 3.23 it was in the $\mathrm{gr}_{\mathbf{w}}$-grading. The basepoint configuration has also changed, however.

Proof. Let us first consider $T^-_{z,w}S^+_{w,z'}$. By Lemma 3.23, we have that

$$S^+_{w,z}\phi_{z\to z'} \simeq S^+_{z',w}.$$

We now just compose with $T^-_{w,z}$ to get

$$T^-_{w,z}S^+_{w,z}\phi_{z\to z'} \simeq T^-_{w,z}S^+_{z',w}.$$

Noting that $T^-_{w,z}S^+_{w,z} \simeq \mathrm{id}$ by Lemma 3.13, the first equivalence now follows. The second equality follows similarly, using Lemma 3.25 instead. □

The above lemma also holds in the case that the basepoints z, w and z' appear in the opposite order, as long as we replace the maps $T^+_{w,z}$ with $T^+_{z,w}$ and $S^-_{z',w}$ with $S^-_{w,z'}$. Analogously, can consider moving a $\mathbf{w}$ basepoint. Suppose $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link with $w \in \mathbf{w}$, and let w' and z' be new basepoints in a component of $L \setminus (\mathbf{w} \cup \mathbf{z})$ adjacent to w, such that the order of the basepoints is w, z' then w'. Let $\phi_{w\to w'}$ denote the diffeomorphism induced by moving w to w' and fixing all of the other basepoints. We have the following:

Lemma 3.28. *The map induced by $\phi_{w\to w'}$ takes the form*

$$(\phi_{w\to w'})_* \simeq S^-_{z',w}T^+_{w',z'} \simeq T^-_{z',w}S^+_{w',z'}.$$

Proof. Applying Lemma 3.25 (note that our labeling of the basepoints is slightly different here), shows that

$$T^+_{w',z'} \simeq T^+_{z',w}(\phi_{w\to w'})_*.$$

Postcomposing with $S^-_{z',w}$ yields

$$S^-_{z',w}T^+_{w',z'} \simeq (\phi_{w\to w'})_*,$$

yielding the first equivalence. The second equivalence follows from a similar argument, using Lemma 3.23. □

We now give a new proof of Equation (4), for curved link Floer homology.

Lemma 3.29. *Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a link with w and z adjacent basepoints in $\mathbf{w}$ and $\mathbf{z}$, such that w follows z. Suppose that w' and z' are new adjacent basepoints which are in the components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ immediately following w. Let ϱ_0 denote the diffeomorphism obtained by moving the pair (w, z) to (w', z'), while fixing all the other basepoints. If z has the same color as the other $\mathbf{z}$ basepoint adjacent to w, then*

$$(\varrho_0)_* \simeq S^-_{w,z}\Psi_{z'}S_{w',z'}.$$

If w has the same color as the other $\mathbf{w}$ basepoint adjacent to z, then

$$(\varrho_0)_* \simeq T^-_{w,z}\Phi_w T^+_{w,z}.$$

For any coloring, we have

$$(\varrho_0)_* \simeq S^-_{w,z}T^+_{w',z'} \simeq T^-_{w,z}S^+_{w',z'}.$$

Proof. We first consider the case of arbitrary colorings. We compute, using our previous computations of the basepoint moving maps in this section

$$\begin{aligned}(\varrho_0)_* &\simeq (\phi_{z\to z'})_*(\phi_{w\to w'})_* \\ &\simeq S^-_{w,z}T^+_{z',w}(\phi_{w\to w'})_* && \text{(Lemma 3.27)} \\ &\simeq S^-_{w,z}T^+_{w',z'} && \text{(Lemma 3.25).}\end{aligned}$$

Analogously, we can also write

$$\begin{aligned}(\varrho_0)_* &\simeq (\phi_{z\to z'})_*(\phi_{w\to w'})_* \\ &\simeq T^-_{w,z}S^+_{z',w}(\phi_{w\to w'})_* && \text{(Lemma 3.27)} \\ &\simeq T^-_{w,z}S^+_{w',z'} && \text{(Lemma 3.23).}\end{aligned}$$

We now consider the case where we have an extra assumption on the colorings of the basepoints. The extra assumptions on the colorings of the basepoints are necessary so that the maps Ψ_z or Φ_w can be defined. With the appropriate assumptions on the colorings, we have by Lemmas 3.27 and 3.28 that

$$\begin{aligned}(\varrho_0)_* &\simeq (\phi_{z\to z'})_*(\phi_{w\to w'})_* \\ &\simeq S^-_{w,z}T^+_{z',w}T^-_{z',w}S^+_{w',z'} && \text{(Lemmas 3.27 and 3.28)} \\ &\simeq S^-_{w,z}\Psi_{z'}S^+_{w',z'} && \text{(Lemma 3.12)}\end{aligned}$$

and similarly

$$\begin{aligned}(\varrho_0)_* &\simeq (\phi_{z\to z'})_*(\phi_{w\to w'})_* \\ &\simeq T^-_{w,z}S^+_{z',w}S^-_{z',w}T^+_{w',z'} && \text{(Lemmas 3.27 and 3.28)} \\ &\simeq T^-_{w,z}\Phi_w T^+_{w',z'} && \text{(Lemma 3.11),}\end{aligned}$$

completing the proof. □

3.7. **Disk-stabilization maps $S^{\pm}_{\mathbb{U},D}$.** Suppose that $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is a multibased link in Y^3 and that $\mathbb{U} = (U, w, z)$ is an unknot which is unknotted from all of the components L. Suppose in addition that D is a choice of embedded disk in $Y \setminus L$ which has boundary U. Write $\mathbb{L} \cup \mathbb{U}$ for the link $(L \cup U, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\})$. Given a coloring σ of $\mathbb{L}$ which is extend by a coloring σ' of $\mathbb{L} \cup \mathbb{U}$, we will now define maps

$$S^+_{\mathbb{U},D} : \mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, \mathbb{L} \cup \mathbb{U}, \sigma', \mathfrak{P}, \mathfrak{s})$$

and

$$S^-_{\mathbb{U},D} : \mathcal{C}FL^\circ(Y, \mathbb{L} \cup \mathbb{U}, \sigma', \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$$

which are well defined up to filtered equivariant chain homotopy. Formally these are constructed very similarly to the free stabilization maps defined in [Zem15b]. The Floer homological component of the proof of invariance of these maps is identical to the proof of invariance of the free-stabilization maps in [Zem15b, Section 6], so we will focus on the topology in this section.

We pick a regular neighborhood $\overline{D}$ of D which doesn't intersect L and define maps $S^{\pm}_{\mathbb{U},D,\overline{D}}$ depending on the choice of regular neighborhood $\overline{D}$. We consider only diagrams $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ such that

(1) $\overline{D} \cap \Sigma$ is a disk;
(2) $\overline{D} \cap (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z}) = \varnothing$.

We call such diagrams **subordinate** to $\overline{D}$. As in [Zem16] and [Zem15b] we also need to choose a collection of data $\mathcal{J}$, which we call gluing data. The collection $\mathcal{J}$ consists of a choice of almost complex structures on $\mathcal{H}$, a choice of embedding of the diagram shown in Figure 3.5 into Σ, as well as some additional data.

Given such a subordinate diagram $\mathcal{H}$ and a choice of gluing data $\mathcal{J}$, we define a diagram $\overline{\mathcal{H}}_{\overline{D},\mathcal{J}}$ by simply replacing $\Sigma \cap \overline{D}$ with a disk shown in Figure 3.5, according to an embedding specified by $\mathcal{J}$. For a given choice of neck length T, a choice of gluing data $\mathcal{J}$ determines an almost complex structure $\mathcal{J}(T)$ on the diagram $\overline{\mathcal{H}}_{\overline{D},\mathcal{J}}$.

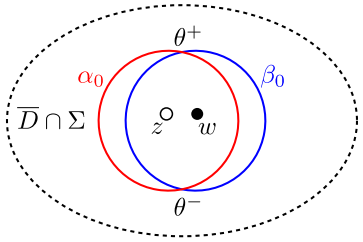

FIGURE 3.5. The diagram we insert into $\overline{D} \cap \Sigma$ for the maps $S^{\pm}_{\mathbb{U},D,\overline{D}}$.

The two intersection points of α_0 and β_0 are distinguished by the relative grading, and the designation of top and bottom generators is the same for $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$. With respect to these diagrams and almost complex structures, we have the following differential computation:

Lemma 3.30. *Suppose that $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a Heegaard diagram for $\mathbb{L}$ and that $\overline{D} \cap \Sigma$ is contained in the component A of $\Sigma \setminus \boldsymbol{\alpha}$ containing, and the component B of $\Sigma \setminus \boldsymbol{\beta}$. Let w_A and z_A (resp. w_B and z_B) denote the basepoints in A (resp. B). Then for sufficiently large T, there is an identification of*

$$\partial_{\mathcal{H},\mathcal{J}(T)} = \begin{pmatrix} \partial_{\mathcal{H}} & * \\ 0 & \partial_{\mathcal{H}} \end{pmatrix}.$$

For a choice of almost complex structure with connected sum point sufficiently close to the α_0 curve, there is an identification of $*$ as $U_w V_z + U_{w_A} V_{z_A}$. For connected sum point sufficient close to the β_0 curve, there is an identification of $*$ as $U_w V_z + U_{w_B} V_{z_B}$.

Proof. The proof is identical to the proof of [OS08, Prop. 6.5]. Alternatively, all components except the top right are computed in [Zem15b, Lem. 6.5], though the proof in the latter is essentially identical to the former argument. □

We define a map

$$S^+_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T} : \mathcal{CFL}^\circ_{J_s}(\mathcal{H}, \sigma_1, \mathfrak{P}, \mathfrak{s}) \to \mathcal{CFL}^\circ_{\mathcal{J}(T)}(\overline{\mathcal{H}}_{\overline{D}}, \sigma_2, \mathfrak{P}, \mathfrak{s}),$$

and a map $S^-_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}$, defined in the opposite direction, by the familiar formulas of

$$S^+_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}(\mathbf{x}) = \mathbf{x} \times \theta^+$$

and

$$S^-_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}(\mathbf{x} \times \theta^-) = \mathbf{x} \qquad \text{and} \qquad S^-_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}(\mathbf{x} \times \theta^+) = 0.$$

For large enough T, the maps $S^\pm_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}$ are chain maps by Lemma 3.30, and also for all sufficiently large T, they commute with increasing the parameter T. As with the free-stabilization and quasi-stabilization maps, we can show that $S^\pm_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}$ are independent of T and $\mathcal{J}$ by a computation of the change of almost complex structure map. Independence from $\mathcal{H}$ follows identically to the free-stabilization maps. Independence from the choice of regular neighborhood $\overline{D}$ follows since any two choices of $\overline{D}$ are isotopic via an isotopy fixing D and L, and the relevant diagrams tautologically commute because the change of diagram maps are induced by isotopies. We have the following:

Proposition 3.31. *For fixed D, the maps $S^\pm_{\mathbb{U},D,\overline{D},\mathcal{H},\mathcal{J},T}$ are independent of the choices of $\overline{D}, \mathcal{H}, \mathcal{J}$ and T.*

Proof. The proof follows identically to [Zem15b, Theorem 6.22]. A key component of the proof is a triangle count in [Zem15b, Theorem 6.17] for freely stabilized triangles. □

We denote the common map by $S^\pm_{\mathbb{U},D}$. We make no claim about independence from the disk D, and indeed we expect the map to depend on the choice of embedded disk D.

4. Maps for handles attached away from L

In this section we define maps associated to handle attachments away from the link L. We define maps associated to surgery on a framed 0-sphere, a framed 1-dimensional link, or a framed 2-sphere. The 1-handle and 3-handle maps are formally identical to the maps defined in [Zem15b] (which are, of course, modifications of the original maps defined in [OS06]). The 2-handle maps are defined nearly identically to the maps from [OS06], though there is a minor difference as in our setting one must pick either the $\mathbf{w}$-basepoints or the $\mathbf{z}$-basepoints to compute Spin^c structures. We use the $\mathbf{w}$-basepoints for this.

Adapting the terminology of [Juh15b], we make the following definition:

Definition 4.1. *We say $\mathbb{S}$ is a **framed k-sphere** in Y, away from L, if $\mathbb{S}$ is an embedding of $S^k \times D^{3-k}$ into Y.*

We usually write $\iota : S^k \times D^{3-k} \to Y$ for the actual map.

4.1. **Model link cobordisms for handles away from L.** We now describe model cobordisms, corresponding to handle attachments away from L. Suppose that L is a multi-based link in Y, and $\mathbb{S}$ is collection of pairwise disjoint, framed spheres in Y, all away from L. If $\mathbb{S}$ has only a single component, we define

$$Y(\mathbb{S}) = (Y \setminus \iota(S^k \times \operatorname{int} D^{3-k})) \cup_\iota (D^{k+1} \times S^{3-k-1}).$$

If $\mathbb{S}$ has multiple components, we iterate this procedure for each component. We allow the possibility that $\mathbb{S} = \varnothing$.

We define the model cobordism $W(Y, L, \mathbb{S})$ for attaching handles along $\mathbb{S}$ by $W(Y, L, \mathbb{S}) = (W, \Sigma)$ where

$$W = Y \times [0,1] \cup_{\mathbb{S}} H, \qquad \text{and} \qquad \Sigma = L \times [0,1].$$

Here H is a disjoint union consisting of one copy of $D^{k+1} \times D^{3-k}$ for each framed k-sphere of $\mathbb{S}$.

4.2. **Maps for framed 0- or 2-spheres.** For $(Y, \mathbb{L})$ a closed 3-manifold with embedded multibased link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, we now define maps associated to framed 0-spheres $\mathbb{S}_0 \subseteq Y \setminus L$ and 2-spheres $\mathbb{S}_2 \subseteq Y \setminus L$. Given a 0-sphere $\mathbb{S}_0 = \{p_1, p_2\} \times D^3$, and a Spin^c structure $\mathfrak{s}$ on Y, there is a unique Spin^c structure $\overline{\mathfrak{s}}$ on the cobordism $W(Y, \mathbb{S}_0)$ which extends $\mathfrak{s}$. We let $\mathfrak{s}'$ denote the restriction to $Y(\mathbb{S}_0)$.

The 1-handle maps will take the form

$$F_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}}} : \mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y(\mathbb{S}_0), \mathbb{L}, \mathfrak{s}').$$

Given a framed 2-sphere $\mathbb{S}_2 \subseteq Y$, there is a canonical "dual" framed zero sphere $\mathbb{S}_0 \subseteq Y(\mathbb{S}_2)$. If $\iota : S^2 \times D^1 \to Y$ is the map associated to $\mathbb{S}_2$, then by definition

$$Y(\mathbb{S}_2) = (Y \setminus \iota(S^2 \times D^1)) \cup_\iota (D^3 \times S^0).$$

The subset $D^3 \times S^0$ is the canonical dual framed zero sphere, $\mathbb{S}_0$.

Since $\mathbb{S}_0$ is a framed sphere, i.e. an embedding of $S^0 \times D^3$ into Y, the manifold $Y(\mathbb{S}_2)(\mathbb{S}_0)$ is canonically diffeomorphic to Y. Thus, to define maps for framed 2-spheres, it is sufficient to instead define maps

$$G_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}}} : \mathcal{C}FL^\circ(Y(\mathbb{S}_0), \mathbb{L}, \mathfrak{s}') \to \mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}),$$

for $\overline{\mathfrak{s}}$ the unique Spin^c structure on $W(Y(\mathbb{S}_0), \mathbb{S}_2)$ extending $\mathfrak{s}'$ on $Y(\mathbb{S}_0)$.

We will only consider diagrams $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for $(Y, \mathbb{L})$ which are **subordinate** to the framed 0-sphere $\mathbb{S}_0$, i.e. $\mathbb{S}_0 \cap \Sigma$ consists of two disks which don't intersect $(\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$. For a choice of data $\mathcal{J}$, which as usual we call gluing data, we define a diagram $\mathcal{H}_\mathcal{J} = (\overline{\Sigma}, \overline{\boldsymbol{\alpha}}, \overline{\boldsymbol{\beta}}, \mathbf{w}, \mathbf{z})$, which is formed by removing small disks around the points p_1 and p_2 and connecting the two new boundary components with an annulus. On the annulus we add two homotopically nontrivial curves, α_0 and β_0, which intersect each other transversely at two points. The two intersection points are distinguished by the relative grading (note that both $\mathrm{gr}_{\mathbf{w}}$ or $\mathrm{gr}_{\mathbf{z}}$ give the same decomposition). The 1- and 3- handle maps take the familiar form:

$$F_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}},\mathcal{H},\mathcal{J}}(\mathbf{x}) = \mathbf{x} \times \theta^+$$

and

$$G_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}},\mathcal{H},\mathcal{J}}(\mathbf{x} \times \theta^+) = 0, \qquad \text{and} \qquad G_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}},\mathcal{H},\mathcal{J}}(\mathbf{x} \times \theta^-) = \mathbf{x}.$$

This gives a 1-handle map from $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$ to $\mathcal{C}FL^\circ(\overline{\mathcal{H}}, \mathfrak{s}')$ and a 3-handle map from $\mathcal{C}FL^\circ(\overline{\mathcal{H}}, \mathfrak{s}')$ to $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$. For a choice of $\mathcal{H}$ and $\mathcal{J}$, by pre- and post-composing with change of diagrams maps, we get a well defined map

$$F_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}},\mathcal{H},\mathcal{J}} : \mathcal{C}FL^\circ(Y, \mathbb{L}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y(\mathbb{S}_0), \mathbb{L}, \mathfrak{s}')$$

and a map $G_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}}}$ in the opposite direction, both of which are maps of coherent filtered chain homotopy types. The above maps *a-priori* depend on the choice of $\mathcal{H}$ and $\mathcal{J}$, though we have the following:

Lemma 4.2. *The maps $F_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}},\mathcal{H},\mathcal{J}}$ and $G_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}},\mathcal{H},\mathcal{J}}$ are independent of $\mathcal{H}$ and $\mathcal{J}$.*

The proof in [Zem15b, Sec. 8] carries over without change.

Lemma 4.3. *The maps $S^\pm_{w,z}$ and $T^\pm_{w,z}$ commute with the 1- and 3-handle maps $F_{Y,\mathbb{S}_0,\overline{\mathfrak{s}}}$ and $F_{Y,\mathbb{S}_2,\overline{\mathfrak{s}}}$.*

Proof. The formulas for the maps obviously commute, but it is not clear whether a single almost complex structure can be used for both maps. However both maps are defined for particular choices of almost complex structures, and it is not clear whether the same structure can be used for both. Instead we will show that the change of almost complex structure map to go between almost complex structures which can be used to compute the respective maps preserves the formulas for the maps.

Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ be a diagram for $\mathbb{L}$, the link in the 3-manifold Y before surgery on $\mathbb{S}_0$ or quasi-stabilization. Suppose that $\mathbb{S}_0 = \{p_1, p_2\} \times D^3$ is a framed 0-sphere, $p \in \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \mathbf{w} \cup \mathbf{z})$ is a point at which we will perform quasi-stabilization. Let α_0 and β_0 denote the new curves from the 1-handle operation. Let α_0' and β_0' denote the two new curves from quasi-stabilization. These are shown in Figure 4.1.

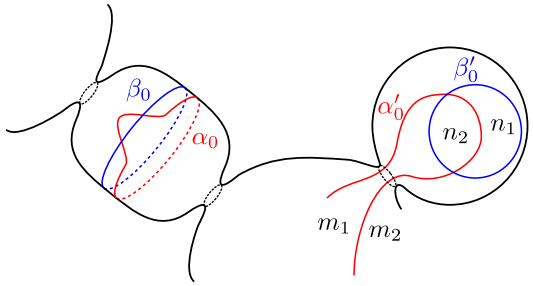

FIGURE 4.1. The three necks we stretch in Lemma 4.3. Multiplicities m_1, m_2, n_1, and n_2 are also labeled.

To show that the 1- and 3-handle maps, as well as the quasi-stabilization maps can all be computed simultaneously we will show that any nonconstant disks counted by the change of almost complex structures map

$$\Phi_{\mathcal{J}(T_1,T_2,T_3)\to\mathcal{J}(T_1',T_2',T_3')},$$

occur for disks

$$\phi \in \pi_2(\mathbf{x} \times x \times x', \mathbf{y} \times y \times y')$$

with $\Delta(x,y) < 0$, where $\Delta(x,y) = \Delta_{\mathbf{w}}(x,y) = \Delta_{\mathbf{z}}(x,y)$ denotes the drop in grading from x to y. Here $x, y \in \alpha_0 \cap \beta_0$ correspond to the intersection points resulting from the 1-handle addition, and $x', y' \in \alpha_0' \cap \beta_0'$ are the intersection points from quasi-stabilization. It is easy to check that this is algebraically sufficient to imply that the quasi-stabilization maps commute with the 1- and 3-handle maps.

Suppose that u_i is a sequence of Maslov index zero holomorphic disks with dynamic almost complex structure representing a homology class ϕ. As we let $i \to \infty$, we can extract a weak limit to a collection of curves U_Σ on $(\Sigma, \boldsymbol{\alpha}^+, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ and U_S on $(S^2, \alpha_0', \beta_0')$. As in [Zem16, Prop. 5.3], the curves in U_Σ consist exactly of a broken holomorphic curve U_Σ' on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ and boundary degenerations $\mathcal{A}$ on $(\Sigma, \boldsymbol{\alpha}^+)$.

Using the Maslov index computations from [Zem16, Lem. 5.3] and [Zem15b, Lem. 6.3], we see that

$$\mu(\phi) = \mu(U_\Sigma) + \Delta(x,y) + n_1(\phi) + n_2(\phi) + m_1(\mathcal{A}) + m_2(\mathcal{A}) + 2 \sum_{\substack{w_0 \in \mathbf{w} \\ w_0 \neq w'}} n_{w_0}(\phi),$$

where n_1, n_2, m_1, and m_2 are the multiplicities shown in Figure 4.1, and w' is the $\mathbf{w}$-basepoint adjacent to z (the new $\mathbf{z}$-basepoint from quasi-stabilization). If $\Delta(x,y) \geq 0$, then all of the above terms must be zero for $\mu(\phi) = 0$. In particular we conclude that $\mu(U_\Sigma) = 0$, so U_Σ must be a constant disk and $\mathcal{A} = \varnothing$. Note that if $\Delta(x,y) = 1$, no disks are counted.

As we mentioned before, with this, it's relatively easy to check all of the possible combinations of $S^{\pm}_{w,z}$ and $F_{Y,\mathbb{S}_0,\bar{\mathfrak{s}}}$ and $G_{Y(\mathbb{S}_0),\mathbb{S}_0,\bar{\mathfrak{s}}}$ to see that they commute. □

Lemma 4.4. *Suppose that $\mathbb{S}$ and $\mathbb{S}'$ are two framed spheres, which are either 0-dimensional or 2-dimensional. If $\mathfrak{s} \in \mathrm{Spin}^c(W(Y, \mathbb{S}, \mathbb{S}'))$, then*

$$F_{Y(\mathbb{S}),\mathbb{L},\mathbb{S}',\mathfrak{s}|_{W(Y(\mathbb{S}),\mathbb{S}')}} F_{Y,\mathbb{L},\mathbb{S},\mathfrak{s}|_{W(Y,\mathbb{S})}} \simeq F_{Y(\mathbb{S}'),\mathbb{L},\mathbb{S},\mathfrak{s}|_{W(Y(\mathbb{S}'),\mathbb{S})}} F_{Y,\mathbb{L},\mathbb{S}',\mathfrak{s}|_{W(Y,\mathbb{S}')}}.$$

Proof. The formulas obviously commute, though one needs to pay attention to the almost complex structures. The appropriate almost complex structure map can be computed using a modification of the argument in Lemma 4.3. In general, the change of almost complex structure map may have nontrivial entries, but using the analogous computation as in Lemma 4.3, any component which counts nonconstant and nonzero disks must have at least one $\Delta(x_i, y_i)$ negative. It is easy to verify algebraically that such disks do not affect the commutator of the two 1- or 3-handle attachment maps. □

4.3. **Maps for framed 1-dimensional links.** Given a multibased link $\mathbb{L}$ in Y with $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and a framed 1-dimensional link $\mathbb{S}_1 \subseteq Y \setminus L$, we now describe maps $F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ associated to a framed link $\mathbb{S}_1$ and the Spin^c structure $\mathfrak{s} \in \text{Spin}^c(W(Y, \mathbb{S}_1))$.

Let $\mathbb{L}_\alpha$ denote the arcs of L which go from $\mathbf{w}$-basepoints to $\mathbf{z}$-basepoints. Let $\mathbb{L}_\beta$ denote the arcs of L which go from $\mathbf{z}$-basepoints to $\mathbf{w}$-basepoints.

Definition 4.5. *An α-**bouquet**, $\mathcal{B}^\alpha$, for a framed link $\mathbb{S}_1$ in $Y \setminus L$ is a collection of arcs which connect links in $\mathbb{S}_1$ to the arcs in $\mathbb{L}_\alpha$. We assume that there is exactly one arc per component of $\mathbb{S}_1$, and each arc has one endpoint on $\mathbb{S}_1$ and one endpoint in $\mathbb{L}_\alpha$. A β-**bouquet**, $\mathcal{B}^\beta$, for $\mathbb{S}_1$ is defined analogously, but with endpoints on $\mathbb{L}_\beta$ instead.*

Given a bouquet $\mathcal{B}^\alpha$ or $\mathcal{B}^\beta$ for a framed link $\mathbb{S}_1 \subseteq Y \setminus L$, we can consider the sutured manifolds $Y(\mathcal{B}^\alpha)$ and $Y(\mathcal{B}^\beta)$ obtained by removing regular neighborhoods of $L \cup \mathcal{B}^\alpha$ and $L \cup \mathcal{B}^\beta$, and adding sutures corresponding to the basepoints $\mathbf{w}$ and $\mathbf{z}$. We have the following definition:

Definition 4.6. *We say that that a triple $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is **subordinate** to the β-bouquet $\mathcal{B}^\beta$ for a framed 1-dimensional link $\mathbb{S}_1 \subseteq Y \setminus L$ if*

1. *$(\Sigma, \alpha_1, \dots, \alpha_n, \beta_{k+1}, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ yields a sutured diagram for $Y(\mathcal{B}^\beta)$;*
2. *The curves $\beta_1, \dots, \beta_k$ are each meridians of a different component of $\mathbb{S}_1$ and hence $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, \mathbb{L})$;*
3. *The curves $\gamma_{k+1}, \dots, \gamma_n$ are each small Hamiltonian isotopies of the curves $\beta_{k+1}, \dots, \beta_n$ respectively, with $|\beta_i \cap \gamma_j| = 2\delta_{ij}$;*
4. *for $i = 1, \dots, k$, the curve γ_i is induced by the framing of the corresponding link component that β_i is a meridian of.*

If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is subordinate to a β-bouquet, then $(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ represents an unlink in $(S^1 \times S^2)^{\#(g-k)}$ with exactly two basepoints per component. As such, regardless of the coloring, $\mathcal{HFL}^-(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}_0)$ admits a generator $\Theta_{\beta\gamma}$ of the top degree subset by Lemma 3.1.

The 4-manifold $X_{\alpha\beta\gamma}$ becomes diffeomorphic to $W(Y, \mathbb{S}_1)$ after filling in $Y_{\beta\gamma}$ with 3- and 4-handles (see [OS06, Proposition 4.3]). As such, given a $\mathfrak{s} \in \text{Spin}^c(W(Y, \mathbb{S}_1))$, we naturally can restrict an $\mathfrak{s} \in \text{Spin}^c(W(Y, \mathbb{S}_1))$ to $\text{Spin}^c(X_{\alpha\beta\gamma})$. Hence given an $\mathfrak{s} \in \text{Spin}^c(W(Y, \mathbb{S}_1))$, we can define the β-subordinate 2-handle map

$$F^\beta_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s},\mathcal{B}^\beta} : \mathcal{CFL}^\circ(Y, \mathbb{L}, \mathfrak{s}|_Y) \to \mathcal{CFL}^\circ(Y(\mathbb{S}_1), \mathbb{L}, \mathfrak{s}|_{Y(\mathbb{S}_1)})$$

by the formula

$$\mathbf{x} \mapsto F_{\alpha\beta\gamma,\mathfrak{s}}(\mathbf{x} \otimes \Theta_{\beta\gamma}) = \sum_{\substack{\psi \in \pi_2(\mathbf{x}, \Theta_{\beta\gamma}, \mathbf{y}) \\ \mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}}} \#\mathcal{M}(\psi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} V_{\mathbf{z}}^{n_{\mathbf{z}}(\phi)} \cdot \mathbf{y}.$$

The proofs in [OS06], [Juh15a], or [Zem15b] show that these don't depend on the choice of bouquet or Heegaard triple subordinate to them.

One similarly defines what it means for a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ to be subordinate to an α-bouquet. The complex $\mathcal{CFL}^-(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathfrak{s}_0)$ will have a generator $\Theta_{\alpha\beta}$ of the top degree, and we define the α-subordinate 2-handle map simply by

$$F^\alpha_{Y,\mathbb{S}_1,\mathfrak{s},\mathcal{B}^\alpha}(\mathbf{x}) = F_{\alpha\beta\gamma,\mathfrak{s}}(\Theta_{\alpha\beta} \otimes \mathbf{x}).$$

As before the original argument from [OS06] goes through essentially without change, and we see that $F^\alpha_{Y,\mathbb{S}_1,\mathfrak{s},\mathcal{B}^\alpha}$ is independent of the choice of α-bouquet or triple subordinate to it.

Remark 4.7. In [OS06], Ozsváth and Szabó construct a map $\mathfrak{s}_{\mathbf{w}} : \pi_2(\mathbf{x},\mathbf{y},\mathbf{z}) \to \mathrm{Spin}^c(X_{\alpha\beta\gamma})$. Analogously to the maps $\mathfrak{s}_{\mathbf{w}} : \mathbb{T}_\alpha \cap \mathbb{T}_\beta \to \mathrm{Spin}^c(Y_{\alpha\beta})$, the map on triangles depends on the choice of basepoints (cf. Lemma 2.8). Indeed the maps $\mathfrak{s}_{\mathbf{w}}$ and $\mathfrak{s}_{\mathbf{z}}$ are not equal, though this asymmetry should be viewed as an essential feature of the theory, as it leads to an Alexander grading change formula, to appear in [Zema]

We now wish to show that $F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ and $F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ are equal. The proof given by Ozsváth and Szabó in [OS06, Lem. 5.2.] for their invariants essentially carries over without change. We state it as a lemma here and outline their argument briefly.

Lemma 4.8. *The maps $F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ and $F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ induce filtered chain homotopy equivalent maps on the coherent filtered chain homotopy invariants.*

Proof sketch. Let $\mathbb{S}_1'$ be an isotopic copy of $\mathbb{S}_1$, in the direction of the framing of $\mathbb{S}_1$. One finds a Heegaard quadruple $(\Sigma,\boldsymbol{\delta},\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{w},\mathbf{z})$ such that $(\Sigma,\boldsymbol{\delta},\boldsymbol{\alpha},\boldsymbol{\beta},\mathbf{w},\mathbf{z})$ is triple subordinate to an α-bouquet of $\mathbb{S}_1'$, and $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{w},\mathbf{z})$ is a triple subordinate to a β-bouquet of $\mathbb{S}_1$. After surgering on $\mathbb{S}_1$, the collection $\mathbb{S}_1'$ becomes a collection of 0-framed unknots. Similarly after surgering on $\mathbb{S}_1'$, the link $\mathbb{S}_1$ becomes a collection of 0-framed unknots. The standard associativity relations shows that,

$$F^{\alpha}_{Y(\mathbb{S}_1),\mathbb{L},\mathbb{S}_1',\mathfrak{s}_0} \circ F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F^{\beta}_{Y(\mathbb{S}_1'),\mathbb{L},\mathbb{S}_1,\mathfrak{s}_0} \circ F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1',\mathfrak{s}}$$

for appropriate choices of Spin^c structures on the two triangle maps. The second map in each of the compositions can be canceled by a composition of 3-handle maps, implying that $F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1}$ and $F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1'}$ are related by post-composition by a diffeomorphism. □

Since the triangle maps are graded over 4-dimensional Spin^c structures, and there isn't a canonical way to compose Spin^c structures, they instead satisfy a Spin^c composition law. We have the following:

Proposition 4.9. *If $\mathbb{S}_1^1$ and $\mathbb{S}_1^2$ are two disjoint framed 1-dimensional links in Y, then*

$$F_{Y(\mathbb{S}_1^1),\mathbb{L},\mathbb{S}_1^2,\mathfrak{s}_2} \circ F_{Y,\mathbb{L},\mathbb{S}_1^1,\mathfrak{s}_1} \simeq \sum_{\substack{\mathfrak{s}\in\mathrm{Spin}^c(W(Y,\mathbb{S}_1^1\cup\mathbb{S}_1^2)) \\ \mathfrak{s}|_{W(Y,\mathbb{S}_1^1)}=\mathfrak{s}_1,\mathfrak{s}|_{W(Y(\mathbb{S}_1^1),\mathbb{S}_1^2)}=\mathfrak{s}_2}} F_{Y,\mathbb{L},\mathbb{S}_1^1\cup\mathbb{S}_1^2,\mathfrak{s}}.$$

Proof. This follows from associativity of the the triangle maps [OS04a, Theorem 8.16]. □

5. Quasi-stabilization and triangle maps

In [Zem16], the quasi-stabilization operation was described in order to define the maps $S^{\pm}_{w,z}$ and $S^{\pm}_{z,w}$. There were three main holomorphic curve counts which were necessary to show naturality of the maps $S^{\pm}_{w,z}$, namely a count of disks necessary to show that the maps $S^{\pm}_{w,z}$ were chain maps, and two counts of holomorphic triangles necessary to show that the maps commuted with the change of diagrams maps induced by moves of the $\boldsymbol{\alpha}$- and $\boldsymbol{\beta}$-curves. In this section, we review the counts of holomorphic triangles in quasi-stabilized diagrams, and we also modify the holomorphic triangle counts from [Zem16, Theorem 6.5], to get a count of holomorphic triangles which will allow us to compute the relations between the band maps and the quasi-stabilization maps. Later, in Remark 7.14, we sketch a strategy which one could use to construct the quasi-stabilization maps without these computations if they were willing to define the quasi-stabilization maps instead as a composition of a disk stabilization map, a band map, and a diffeomorphism map.

All of statements about higher and lower graded generators in this section will be with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading. Symmetrically, all statements also hold if one replaces $\mathrm{gr}_{\mathbf{w}}$ with $\mathrm{gr}_{\mathbf{z}}$.

5.1. Strongly positive diagrams for $(S^1\times S^2)^{\#k}$ and quasi-stabilization of Heegaard triples.

In this subsection we review the results from [Zem16] about quasi-stabilization and triangle maps. We first review a result from [MO10], which will be useful in showing that quasi-stabilization commutes with the change of diagram maps induced by $\boldsymbol{\beta}$-moves.

Given a Heegaard triple $(\Sigma,\boldsymbol{\alpha},\boldsymbol{\beta},\boldsymbol{\gamma},\mathbf{w},\mathbf{z})$, we quasi-stabilize it along a single α_s curve passing through a point $p\in\Sigma$, as in Figure 5.1. We have the following:

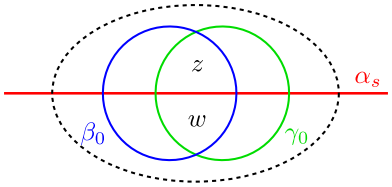

FIGURE 5.1. The version of quasi-stabilization discussed in Lemma 5.1, which is considered in [MO10, Prop. 5.2]

Lemma 5.1 ([MO10, Prop. 5.2]). *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a strongly $\mathfrak{s}$-admissible triple and suppose that α_s is a new $\boldsymbol{\alpha}$-curve, passing through the point $p \in \Sigma$. Let $\overline{\mathcal{T}}_{\alpha_s,p}$ denote the Heegaard triple resulting from quasi-stabilizing along α_s at p, as in Figure 5.1. If $\mathcal{J}$ is gluing data, then for sufficiently large T, with the almost complex structure $\mathcal{J}(T)$ we have the following identifications:*

$$F_{\overline{\mathcal{T}}_{\alpha_s,p},\bar{\mathfrak{s}}}(\mathbf{x}\times\cdot,\mathbf{y}\times y^+) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x},\mathbf{y}) & 0 \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x},\mathbf{y}) \end{pmatrix}$$

where $\cdot \in \{x^+, x^-\} = \alpha_s \cap \beta_0$ and the matrix on the right denotes the expansion into the upper and lower generator components of $\alpha_s \cap \beta_0$ and $\alpha_s \cap \gamma_0$.

We also need an count which shows that the map $S^{\pm}_{w,z}$ is independent of the choice of curve α_s. This requires a more subtle result about how curves can degenerate. Indeed taking an arbitrary Heegaard triple and performing quasi-stabilization along two curves α_s and β_s which intersect at two points in general will not result in curve counts as in the above theorem. Nevertheless, if we restrict which diagrams can appear as $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, we can achieve such a curve count. We make the following definition:

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(S^1 \times S^2)^{\#k}$ such that

$$|\alpha_i \cap \beta_j| = \begin{cases} 1 \text{ or } 2 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

and that if $\alpha_i \cap \beta_i = \{p_i^-, p_i^+\}$, then p_i^- and p_i^+ differ by Maslov grading one. We do not assume that the α_i curves are small isotopies of the β_j curves. Let $\theta^+ = p_1^+ \times \cdots \times p_{n-1}^+$ denote the top graded (partial) intersection point. Assume that $|\alpha_n \cap \beta_n| = 2$ and write p_n^- and p_n^+ for the two points of $\alpha_n \cap \beta_n$.

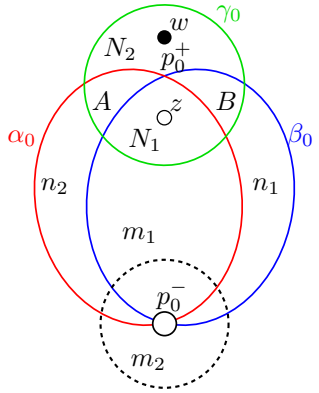

FIGURE 5.2. The diagram we insert into a Heegaard triple diagram $\mathcal{T}$ along the curves α_s and β_s to form the diagram $\overline{\mathcal{T}}_{\alpha_s,\beta_s,p^+}$. We cut out the solid circle marked with p_0^- and stretch the almost complex structure along the dashed circle.

Definition 5.2. *Under the assumptions of the previous paragraph, we say that* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ *is* ***strongly positive*** *with respect to* p_n^+ *if for every nonnegative disk* $\phi \in \pi_2(\theta^+ \times p_n^+, \mathbf{y} \times p_n^+)$ *we have that*

$$(\mu - (m_1 + m_2))(\phi) = (\mu - (n_1 + n_2))(\phi) \geq 0,$$

with equality to zero iff ϕ *is the constant disk. Here* m_1, m_2, n_1, n_2 *denote the multiplicities adjacent to the point* p_n^+*, appearing in the following counterclockwise order:* n_1, m_1, n_2*, then* m_2*.*

Indeed Heegaard triples where $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\})$ is strongly positive with respect to the quasi-stabilization point are ones where we have a reasonable count of curves, as the following theorem demonstrates:

Theorem 5.3 ([Zem16]Theorem 6.5)**.** *Suppose that* $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ *is a Heegaard triple with curves* α_s *and* β_s*, intersecting at two points* p^+ *and* p^- *and let* $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+}$ *be the Heegaard triple result from quasi-stabilization, as described above. If* $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\})$ *is a strongly positive diagram for* $(S^1 \times S^2)^{\#k}$ *with respect to the point* p^+ *(Definition 5.2), and* $\mathcal{J}$ *is gluing data for stretching along the dashed circle, then for sufficiently large* T*, with respect to the almost complex structure* $\mathcal{J}(T)$*, there are identifications*

$$F_{\mathcal{T}^+_{\alpha_s, \beta_s, p^+}, \bar{\mathfrak{s}}, \mathcal{J}(T)}(\Theta^+_{\alpha\beta} \times p_0^+, \mathbf{x} \times \cdot) = \begin{pmatrix} F_{\mathcal{T}, \mathfrak{s}, J_s}(\Theta^+_{\alpha\beta}, \mathbf{x}) & 0 \\ 0 & F_{\mathcal{T}, \mathfrak{s}, J_s}(\Theta^+_{\alpha\beta}, \mathbf{x}) \end{pmatrix},$$

where $\cdot$ *denotes* $x^\pm \in \beta_0 \cap \gamma_0$ *and the matrix on the right denotes the matrix decomposition of the map based on the decompositions given* $x^\pm$ *and* $y^\pm$*. The decompositions are with respect to the* $\mathrm{gr}_{\mathbf{w}}$ *grading.*

5.2. **Weakly positive diagrams and further holomorphic triangle computations.** Unfortunately Theorem 5.3 only covers triangle diagrams which achieve the strong positivity condition. Indeed it does not hold for general triangle diagrams, even for the ones which will occur in the band maps.

Instead, we need a weaker condition. Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ is a Heegaard triple such that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ is a diagram for $(S^1 \times S^2)^{\#k}$ such that each $\alpha_i \cap \beta_i$ consists of two points differing by grading 1, or just one point. Further assume that $\alpha_i \cap \beta_j = \varnothing$ if $i \neq j$. Suppose $\alpha_n \cap \beta_n = \{p_n^-, p_n^+\}$.

Definition 5.4. *Suppose that* $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w})$ *is a triple as in the previous paragraph, and* α_s *and* β_s *are new* $\boldsymbol{\alpha}$ *and* $\boldsymbol{\beta}$ *curves, intersecting each other twice, with top intersection point* p^+*, and let* $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+}$ *denote the quasi-stabilized triple, as in Figure 5.2. We say that* $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+}$ *is* ***weakly positive*** *at* p^+ *if for every disk* $\phi \in \pi_2(\theta^+ \times p^+, \mathbf{y} \times p^+)$ *on* $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\})$ *which is nonnegative, we have that*

$$(\mu - (m_1 + m_2))(\phi) = (\mu - (n_1 + n_2))(\phi) \geq 0,$$

and if we have equality to zero, then $m_2(\phi) = 0$*, where* m_2 *is the multiplicity*[4] *in Figure 5.2.*

Clearly any diagram which is strongly positive at p_n^+ is also weakly positive at p_n^+.

Lemma 5.5. *Suppose that* $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ *is a diagram for* $(S^1 \times S^2)^{\#k}$ *where the curves in* $\boldsymbol{\alpha}$ *are all small Hamiltonian isotopies of the curves in* $\boldsymbol{\beta}$ *with* $\alpha_i \cap \beta_j = 2\delta_{ij}$*. If* $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+} = (\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \boldsymbol{\gamma} \cup \{\gamma_0\}, \mathbf{w} \cup \{w\})$ *is a free-stabilization*[5] *at any point in* $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma} \cup \mathbf{w})$*, then* $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+}$ *is weakly positive at* p^+*.*

Proof. The proof breaks into cases, depending on the region of $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$ where the free-stabilization takes place. If it takes place in a region already containing a basepoint, the diagram $\mathcal{H}$ is actually strongly positive at p^+ by [Zem16, Lemma 6.3], and hence weakly positive as a consequence.

The other alternative is that it occurs in one of the small bigons corresponding to a Maslov index one disk from p_k^+ to p_k^-. Suppose ϕ is a disk in $\pi_2(\theta^+ \times p^+, \mathbf{y} \times p^+)$. Consider the multiplicities $m_1, m_2, n_1, n_2, m_2', n_1'$,

[4]Note that the multiplicity m_2 from Figure 5.2 refers to a multiplicity on $(S^2, \alpha_0, \beta_0, \gamma_0, w)$, the diagram we insert to perform quasi-stabilization, but here we mean the corresponding multiplicity on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w} \cup \{w\})$. Note also the relation of the multiplicity m_2 with the basepoint w.

[5]As in [Zem15b], a free-stabilization is where α_s, β_s and γ_0 are all Hamiltonian isotopic, each pair intersecting in two points, such α_s, β_s, and γ_0 are contained in a ball B on Σ which has trivial intersection with $\boldsymbol{\alpha} \cup \boldsymbol{\beta} \cup \boldsymbol{\gamma} \cup \mathbf{w}$. The basepoint w is added inside the region bounded by all three curves.

and n_2' shown in Figure 5.3. Let $\phi_0 \in \pi_2(\theta^+, \mathbf{y})$ denote the disk induced by ϕ, on the unstabilized diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$.

By assumption $p_k^+ \in \theta^+$. Either p_k^+ or p_k^- is in $\mathbf{y}$. Let $\delta = \delta(p_k^+, \mathbf{y}) \in \{0, 1\}$ be 1 if $p_k^+ \in \mathbf{y}$ and 0, otherwise.

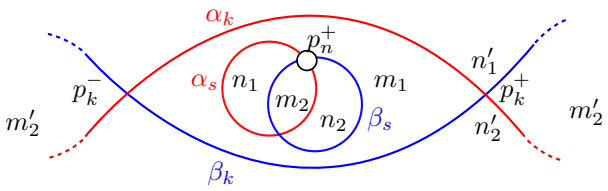

FIGURE 5.3. A free-stabilization inside a bigon. Multiplicities for a disk ϕ are labeled.

By assumption on the diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w})$, the regions with multiplicities n_1' and n_2' must each have a basepoint (though they could be the same region). Hence $n_{\mathbf{w}}(\phi_0) \geq \frac{1}{2}(n_1' + n_2')$. Since $\phi_0 \in \pi_2(\theta^+, \mathbf{y})$, and θ^+ is the top degree intersection point, using the grading formula, we have

$$\mathrm{gr}(\theta^+) - \mathrm{gr}(\mathbf{y}) = \mu(\phi_0) - 2n_{\mathbf{w}}(\phi_0) \geq \delta,$$

which yields

$$\mu(\phi_0) \geq 2n_{\mathbf{w}}(\phi_0) + \delta.$$

Using our bound for $n_{\mathbf{w}}(\phi_0)$, above, we have

$$\mu(\phi_0) \geq n_1' + n_2' + \delta.$$

Using the formula for the Maslov index of a disk after performing free-stabilization (cf. [Zem15b, Lemma 6.3]), we have that

$$\mu(\phi) \geq n_1' + n_2' + 2m_2 + \delta.$$

Since $n_1' + n_2' + \delta = m_1 + m_2'$ we have

$$\mu(\phi) - m_1 - m_2 \geq m_1 + m_2' + 2m_2 - (m_1 + m_2) = m_2' + m_2,$$

which is certainly a nonnegative quantity. On the other hand, if $\mu(\phi) - m_1 - m_2 = 0$, then $m_2' + m_2 = 0$ and hence in particular $m_2 = 0$. □

We now show that weakly positive double quasi-stabilizations have reasonable triangle map counts.

Proposition 5.6. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with curves α_s and β_s, intersecting at two points p^+ and p^- and let $\overline{\mathcal{T}}_{\alpha_s, \beta_s, p^+}$ be the Heegaard triple result from quasi-stabilization, which satisfies the weak positivity condition, with respect to the top generator, p^+, of $\alpha_s \cap \beta_s$. Suppose $\mathcal{J}$ is gluing data for stretching the almost complex structure around γ_0. then for sufficiently large T, with respect to the almost complex structure $\mathcal{J}(T)$, there are identifications*

$$F_{\mathcal{T}^+_{\alpha_s, \beta_s, p^+}, \bar{\mathfrak{s}}, \mathcal{J}(T)}(\Theta^+_{\alpha\beta} \times p_0^+, \mathbf{x} \times \cdot) = \begin{pmatrix} F_{\mathcal{T}, \mathfrak{s}, J_s}(\Theta^+_{\alpha\beta}, \mathbf{x}) & * \\ 0 & F_{\mathcal{T}, \mathfrak{s}, J_s}(\Theta^+_{\alpha\beta}, \mathbf{x}) \end{pmatrix},$$

where $\cdot$ denotes $x^\pm \in \beta_0 \cap \gamma_0$ and the matrix on the right denotes the matrix decomposition of the map based on the decompositions given $x^\pm$ and $y^\pm$. The decomposition is taken with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading.

Proof. We perform the same degeneration as in [Zem16, Theorem 6.5]. Indeed if ψ is a Maslov index zero triangle in $\pi_2(\theta^+ \times p_n^+, \mathbf{x} \times x, \mathbf{y} \times y)$ with holomorphic representatives for arbitrarily large neck lengths, then we extract a weak limit as a collection of curves, as follows:

(1) a broken holomorphic triangle u_Σ representing a homology class ψ_Σ on $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ which has no boundary components on α_s or β_s;
(2) a broken holomorphic disk $u_{\alpha\beta}$ on $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\})$ representing a class $\phi_{\alpha\beta} \in \pi_2(\Theta^+_{\alpha\beta} \times p^+, \mathbf{y} \times p^+)$, for some $\mathbf{y} \in \mathbb{T}_\alpha \cap \mathbb{T}_\gamma$;
(3) a broken holomorphic triangle u_0 on $(S^2, \alpha_0, \beta_0, \gamma_0)$ which represents a homology triangle $\psi_0 \in \pi_2(p_0^+, x, y)$.

As in [Zem16, Theorem 6.5], the Maslov index can be computed to be

$$\mu(\psi) = \mu(\psi_\Sigma) + (N_1 + N_2)(\psi_0) + \mu(\phi_{\alpha\beta}) - \tfrac{1}{2}(n_1 + n_2 + m_1 + m_2)(\phi_{\alpha\beta}).$$

By weak positivity, we know that the sum above can be expressed as a sum of nonnegative integers, and hence each must be zero. That is,

$$\mu(\psi_\Sigma) = N_1(\psi_0) + N_2(\psi_0) = \mu(\phi_{\alpha\beta}) - \tfrac{1}{2}(n_1 + n_2 + m_1 + m_2)(\phi_{\alpha\beta}) = 0.$$

By weak positivity, we also know that $m_2(\phi_{\alpha\beta}) = 0$.

Now since $\mu(\psi_\Sigma) = 0$, we know that u_Σ, the broken triangle representing ψ_Σ, must actually be a genuine holomorphic triangle. As such, there is a genuine holomorphic triangle, u_0^m, in the broken curve u_0, which matches u_Σ (as in the proof of [Zem16]Theorem 6.5). Since $m_2(\phi_{\alpha\beta}) = 0$, we know that $m_2(\psi) = m_2(u_\Sigma) = m_2(u_0^m)$. Now u_0^m is a holomorphic triangle on $(S^2, \alpha_0, \beta_0, \gamma_0)$ which must be in $\pi_2(p_0, x_0, y_0)$, for some choice of intersection points p_0, x_0, and y_0. Now $p_0 = p_0^+$, since

$$m_1(u_0^m) = m_2(u_0^m) = n_1(u_0^m) = n_2(u_0^m)$$

and so p_0^- cannot be a vertex of the triangle u_0^m, which forces p_0^+ to be a vertex. Since u_0^m has vertex p_0^+, and all multiplicities m_1, m_2, n_1 and n_2 equal, it follows that x_0 and y_0 share the same relative grading (i.e. either both the top or both the bottom). We now consider which multiplicities could be increased in going from u_0^m to u_0. The homology class $u_0 - u_0^m$ is some collection of disks, boundary degenerations and closed surfaces. Since

$$N_1(u_0 - u_0^m) = N_2(u_0 - u_0^m) = N_1(\psi_0) = N_2(\psi_0) = 0,$$

it's easy to see that $A(u_0 - u_0^m) = B(u_0 - u_0^m) = 0$ as well. Since $m_2(u_0 - u_0^m) = 0$ as well, we observe that $u_0 - u_0^m$ is constrained to be one of the two Maslov index one disks shown in Figure 5.4.

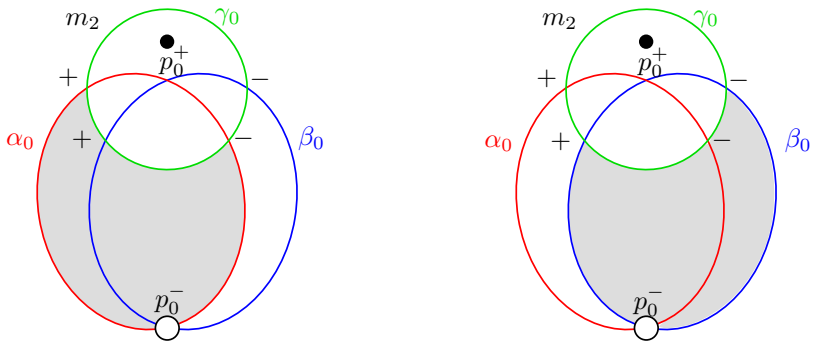

FIGURE 5.4. The two homology disks in Proposition 5.6 which may appear in the homology class of $u_0 - u_0^m$.

The one on the left can be added to u_0 if u_0 represents a class in $\pi_2(p_0^+, x_0^+, y_0^+)$, and the one on the right can be added to u_0 if u_0 represents a class in $\pi_2(p_0^+, x_0^-, y_0^-)$. In either case, the resulting triangle is in $\pi_2(p_0^+, x^-, y^+)$, corresponding to the component marked with an asterisk in the proposition statement.

Since such triangles in $\pi_2(p_0^+, x^-, y^+)$ are excluded from our computation, we conclude that $u_0^m = u_0$. This forces $(m_1 + m_2 + n_1 + n_2)(\phi_{\alpha\beta}) = 0$. As $\mu(\phi_{\alpha\beta}) - \frac{1}{2}(m_1 + m_2 + n_1 + n_2)(\phi_{\alpha\beta}) = 0$, we conclude that $\mu(\phi_{\alpha\beta}) = 0$, and since $\phi_{\alpha\beta}$ admits a broken holomorphic representative, we conclude that $\phi_{\alpha\beta}$ is the constant disk. Hence ψ is represented by exactly a prematched triangle, and ψ is a class in $\pi_2(\theta^+ \times p^+, \mathbf{x} \times x, \mathbf{y} \times y)$ for x and y of the same relative grading. From here the proof follows the proof of [Zem16, Theorem 6.5] identically. □

6. Definition of the band maps

In this section, we construct maps associated to surgering a link $\mathbb{L} \subseteq Y$ along a band B embedded in Y. A band B on a link L is an embedded copy of $[-1,1] \times [-1,1]$ in Y such that $B \cap L$ consists of exactly $\{-1,1\} \times [-1,1]$. In terms of link cobordisms, a band corresponds to a saddle critical point of the surface in

the link cobordism. If L is an oriented link, an **oriented** band B is a band such that the orientation of B is opposite the orientation of L on $B \cap L$. If L is an oriented link and B is an oriented band for L, we can form the surgered link, $L(B)$, which is also an oriented link. We will assume all bands are oriented, and also omit stating that bands are oriented.

In terms of decorated cobordisms, we can assume that if $(W, (\Sigma, \mathcal{A}))$ is a decorated link cobordism, and f is a Morse function on W which restricts to a Morse function on Σ, then generically the index 1 critical points of $f|_\Sigma$ occur away from the divides $\mathcal{A}$. They can occur either inside a type-$\mathbf{w}$ region or inside of a type-$\mathbf{z}$ region. We will define two maps, which we will denote by $F^{\mathbf{z}}_B$ and $F^{\mathbf{w}}_B$ corresponding to bands consisting of a type-$\mathbf{z}$ region or a type-$\mathbf{w}$ region, respectively.

Maps associated to band moves between links have been described (implicitly or explicitly) by many authors. For example, such maps have been defined explicitly by Sarkar in [Sar11b], Ozsváth, Stipsicz and Szabó in [OSS15]. Usually they are defined by a triangle map, or using grid diagrams. Alternatively, they can be defined in the hat flavor of link Floer homology using the construction in [Juh15a], though this approach is somewhat different than the others in that it requires the gluing map in [HKM08] defined using contact geometry. In this section we construct maps associated to bands on $\mathcal{C}FL^\circ$, and prove that the maps are invariant from the choices in the construction.

Additionally, it's worth mentioning that in [JM16], Juhász and Marengon prove a skein exact sequence for the hat version of link Floer homology, using link cobordism maps. In light of this, it is reasonable to suspect that many of the maps fitting into such exact sequences, which are often computed by counting holomorphic triangles, are secretly link cobordism maps. Examples of such constructions can be found in, for example, [OS04b] or [Man07].

6.1. **Topological preliminaries on bands.** We wish to consider bands on oriented, multi-based links, and we now define the type of bands which will be useful for our theory.

Definition 6.1. *We say that a band B^α of an oriented multi-based link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ is an α-**band** if the the ends of B^α occur in components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which go from $\mathbf{w}$-basepoints to $\mathbf{z}$-basepoints. We say that a band B^β is a β-**band** if the ends of B occur in components of $L \setminus (\mathbf{w} \cup \mathbf{z})$ going from $\mathbf{z}$-basepoints to $\mathbf{w}$-basepoints. For a band B of either type, we always assume that the ends of B lie in distinct components of $L \setminus (\mathbf{w} \cup \mathbf{z})$.*

Notice that a band can never be both an α-band and a β-band. If $\mathbb{L}$ is a multibased link and B is a band (of either type) for a link $\mathbb{L}$ in Y, we can form the surgered link $L(B)$ in the obvious way. Since the ends of B are assumed to lie distinct components of $L \setminus (\mathbf{w} \cup \mathbf{z})$, the link with basepoints

$$\mathbb{L}(B) = (L(B), \mathbf{w}, \mathbf{z}),$$

is a multi-based link in the sense of Definition 2.1.

Definition 6.2. *We say that that a Heegaard triple $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is **subordinate** to an α-band B^α for $\mathbb{L}$, if*

1. *the diagram $(\Sigma, \beta_1, \dots, \beta_{n-1}, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ induces a sutured diagram for the sutured manifold $(Y \setminus N(L \cup B), \partial N(\mathbf{w} \cup \mathbf{z}))$, where $\partial N(\mathbf{w} \cup \mathbf{z}) \subseteq \Sigma$ are given orientations compatible to the orientation of L, as in* [Juh15a, Definition 4.4]*;*
2. *$(\Sigma, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, \mathbb{L})$;*
3. *$\alpha_1, \dots, \alpha_{n-1}$ are small Hamiltonian isotopies of the curves $\beta_1, \dots, \beta_{n-1}$ with $|\alpha_i \cap \beta_j| = \delta_{ij}$, for $i, j \in \{1, \dots, n-1\}$;*
4. *the curve α_n is induced by the band B^α and $(\Sigma, \alpha_1, \dots, \alpha_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\alpha))$.*

Using [JT12, Proposition 2.37], we know that any two sutured Heegaard diagrams for $Y \setminus N(L \cup B)$ can be connected by a set of Heegaard moves for embedded Heegaard surfaces, and hence we have the following:

Lemma 6.3. *Given an α-band B^α, any two triples subordinate to B^α can be connected by a sequence of the following moves*

1. *isotopies of Σ inside of Y which fix $L \cup B$.*
2. *isotopies or handleslides amongst the $\boldsymbol{\gamma}$-curves;*

(3) *isotopies or handleslides amongst the $\alpha_1, \dots, \alpha_{n-1}$ curves followed by the corresponding handleslide or isotopy of the corresponding curve in $\beta_1, \dots, \beta_{n-1}$;*
(4) *isotopies or handleslides of the curve β_n across other $\boldsymbol{\beta}$-curves;*
(5) *isotopies or handleslides of the curve γ_n across other $\boldsymbol{\gamma}$-curves;*
(6) *$(1,2)$-stabilizations of the Heegaard triple, i.e., taking the internal connected sum of Σ in Y with a torus $\mathbb{T}^2 \subseteq Y$ such that there is a 3-ball $B^3 \subseteq Y$ containing $\mathbb{T}^2$ with $B^3 \cap (\Sigma \cup L \cup B) = \varnothing$. On $\mathbb{T}^2$ we put three curves, α_0,β_0 and γ_0, with $|\beta_0 \cap \gamma_0| = 1$ and α_0 a small Hamiltonian isotopy of β_0 with $|\alpha_0 \cap \beta_0| = |\alpha_0 \cap \gamma_0| = 2$.*

There is an analogous lemma and definition for triples subordinate to β-bands B:

Definition 6.4. *We say a Heegaard triple $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ is **subordinate** to a β-band B^β if*

(1) *$(\Sigma \setminus N(\mathbf{w} \cup \mathbf{z}), \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_{n-1}, \partial N(\mathbf{w} \cup \mathbf{z}))$ induces a sutured diagram for the sutured manifold $(Y \setminus N(L \cup B^\beta), \partial N(\mathbf{w} \cup \mathbf{z}))$, where the circles $\partial N(\mathbf{w} \cup \mathbf{z})$ are given orientations compatible to the orientation of L, as in* [Juh15a, Definition 4.4]*;*
(2) *$(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ is a diagram for $(Y, \mathbb{L})$;*
(3) *$\gamma_1, \dots, \gamma_{n-1}$ are small Hamiltonian isotopies of the curves $\beta_1, \dots, \beta_{n-1}$, with $|\beta_i \cap \gamma_j| = \delta_{ij}$, for $i, j \in \{1, \dots, n-1\}$;*
(4) *the curve γ_n is induced by the band B^β and $(\Sigma, \alpha_1, \dots, \alpha_n, \gamma_1, \dots, \gamma_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\beta))$.*

As before we have the following:

Lemma 6.5. *Given a β-band B^β, any two triples subordinate to B^β can be connected by a sequence of the following moves*

(1) *isotopies of Σ inside of Y which fix $L \cup B^\beta$.*
(2) *isotopies or handleslides amongst the $\boldsymbol{\alpha}$-curves;*
(3) *isotopies or handleslides amongst the $\beta_1, \dots, \beta_{n-1}$ curves followed by the corresponding handleslide or isotopy of the corresponding curve in $\gamma_1, \dots, \gamma_{n-1}$;*
(4) *isotopies or handleslides of the curve β_n across other $\boldsymbol{\beta}$-curves;*
(5) *isotopies or handleslides of the curve γ_n across other $\boldsymbol{\gamma}$-curves;*
(6) *$(1,2)$-stabilizations of the Heegaard triple, by taking the internal connected sum of Σ in Y with a torus $\mathbb{T}^2 \subseteq Y$ such that there is a 3-ball $B^3 \subseteq Y$ containing $\mathbb{T}^2$ with $B^3 \cap (\Sigma \cup L \cup B^3) = \varnothing$. On $\mathbb{T}^2$ we put three curves, α_0,β_0 and γ_0, with $|\alpha_0 \cap \beta_0| = 1$ and γ_0 a small Hamiltonian isotopy of β_0 with $|\alpha_0 \cap \gamma_0| = |\beta_0 \cap \gamma_0| = 2$.*

We now define maps associated to bands. We first consider α-bands.

Lemma 6.6. *If $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a diagram subordiate to an α-band B, then $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a diagram for an unlink in (a disjoint union of copies of) $(S^1 \times S^2)^{\#k}$, and every component of the unlink has exactly two basepoints, except one component, which has four.*

Proof. Let $(Y_{\alpha\beta}, \mathbb{L}_{\alpha\beta})$ denote the 3-manifold with embedded link obtained by attaching 2-handles to $\Sigma \times [0,1]$ and filling in the resulting boundary spheres in the unique way. The link $\mathbb{L}_{\alpha\beta}$ is obtained by concatenating a simple closed curve in each of the boundary spheres. Note that each pair (α_i, β_i) for $i = 1, \dots, n-1$ determines an embedded 2-sphere S_i in $Y_{\alpha\beta}$ which doesn't intersect $\mathbb{L}_{\alpha\beta}$. By surgering out S_i, we can reduce to the case that we are left with $|L|$ genus zero components of Σ. There are $|L|-1$ of these components which will have no $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ curves, but will have a single pair of $\mathbf{w}$ and $\mathbf{z}$ basepoints (thus determining an unknotted component of $\mathbb{L}_{\alpha\beta}$ in $Y_{\alpha\beta}$ with exactly two basepoints). There will be one component of the resulting Heegaard diagram which contains the curves α_n and β_n. In such a manner we reduce the claim to the case that the band B^α is a band between two unknots in S^3, each with two basepoints. A model diagram shows that for this link, α_n and β_n can be taken intersect twice on a genus zero Heegaard surface such that each of the resulting regions on Σ contains a single basepoint, completing the proof. □

6.2. **Maps for type-z bands.** Given a type $\mathbf{z}$-band B, for a multi-based link $\mathbb{L}$ with coloring $\sigma : (\mathbf{w} \cup \mathbf{z}) \to \mathfrak{P}$, we now define a map

$$F_B^{\mathbf{z}} : \mathcal{CFL}^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{CFL}^\circ(Y, \mathbb{L}(B), \sigma, \mathfrak{P}, \mathfrak{s}).$$

These are defined only when the $\mathbf{z}$-basepoints adjacent to the ends of the band have the same color.

As we showed in Lemma 3.2, there is a distinguished generator as a $\mathbb{Z}_2[V_{\mathbf{z}}]$-module

$$\Theta^{\mathbf{w}}_{\alpha\beta} \in \mathcal{HFL}^-_{\max,\mathbf{w}}(\mathcal{H}_{\alpha\beta}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}_0) \subseteq \mathcal{HFL}^-(\mathcal{H}_{\alpha\beta}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}_0)$$

where $\mathcal{HFL}^-_{\max,\mathbf{w}}$ denotes the subgroup of top homological degree with respect to the grading $\mathrm{gr}_{\mathbf{w}}$, and $(\sigma_0, \mathfrak{P}_0)$ is the coloring that sets V_z equal to $V_{z'}$, as in Lemma 3.2. Note that for any other coloring which gives z and z' the same color, there is a chain map

$$\mathcal{HFL}^-(\mathcal{H}_{\alpha\beta}, \sigma_0, \mathfrak{P}_0, \mathfrak{s}_0) \to \mathcal{HFL}^-(\mathcal{H}_{\alpha\beta}, \sigma, \mathfrak{P}, \mathfrak{s}_0),$$

so the element $\Theta^{\mathbf{w}}_{\alpha\beta}$ is a well defined element of $\mathcal{HFL}^-$ for any coloring that assigns z and z' the same color.

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple subordinate to an α-band. The triangle maps require a choice of Spin^c structure on $X_{\alpha\beta\gamma}$, but we note that $X_{\alpha\beta\gamma}$ becomes diffeomorphic to $Y \times [0,1]$ after filling in $Y_{\alpha\beta}$ with 3- and 4-handles. Letting $\mathrm{Spin}^c(X_{\alpha\beta\gamma})_0 \subseteq \mathrm{Spin}^c(X_{\alpha\beta\gamma})$ denote the set of Spin^c structures which restrict to the torsion Spin^c structure on $Y_{\alpha\beta}$, we have

$$\mathrm{Spin}^c(X_{\alpha\beta\gamma})_0 \cong \mathrm{Spin}^c(Y).$$

Hence given an α-band B^α and a Spin^c structure $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, we can define a band map

$$F^{\mathbf{z}}_{B^\alpha}(\cdot) = F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}(\Theta^{\mathbf{w}}_{\alpha\beta} \otimes \cdot) : \mathcal{CFL}^\circ(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{CFL}^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}).$$

Note that by $F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}$ we mean the map which counts holomorphic triangles representing classes ψ with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s}$. For links where the sum of all of the components are nullhomologous, this presents no asymmetry between the $\mathbf{w}$ and $\mathbf{z}$ basepoints as $\mathfrak{s}_{\mathbf{w}} = \mathfrak{s}_{\mathbf{z}}$ by Lemma 2.8 (or rather, it's analog for Heegaard triples, see [Zema]). For links which aren't nullhomologous in this sense, this introduces an asymmetry between the $\mathbf{w}$ and $\mathbf{z}$ basepoints.

We now claim that these induce well defined maps on the coherent chain homotopy type:

Lemma 6.7. *The map $F^{\mathbf{z}}_{B^\alpha}$ is independent (up to equivariant, filtered chain homotopy) of the choice of Heegaard triple subordinate to a framed arc for the band B^α, and hence descends to a well defined map on the coherent filtered chain homotopy type invariants.*

Proof. We just need to check independence from each of the moves appearing Lemma 6.3. Each of these moves is considered in the maps associated to surgery on framed 1-dimensional links embedded in 3-manifolds, and can be addressed in the same way (cf. [Zem15b, Lemma 9.12] and [Juh15a, Theorem 6.9]). □

Analogously, if $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a triangle subordinate to a β-band B^β, there is an element

$$\Theta^{\mathbf{w}}_{\beta\gamma} \in \mathcal{HFL}^-_{\max,\mathbf{w}}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}_0)$$

which generates $\mathcal{HFL}^-_{\max,\mathbf{w}}$ as a $\mathbb{Z}_2[V_{\mathbf{z}}]$-module. Thus as long as $\mathbf{z}$-basepoints adjacent to the ends of the band B^β are given the same color, we can define a band map

$$F^{\mathbf{z}}_{B^\beta}(\cdot) = F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}(\cdot \otimes \Theta^{\mathbf{w}}_{\beta\gamma}) : \mathcal{CFL}^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{CFL}^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}),$$

which descends to the coherent $\mathfrak{P}$-filtered chain homotopy type invariant.

Remark 6.8. In [OS06], Ozsváth and Szabó define cobordism maps for performing surgery on framed links by taking Heegaard triples which are subordinate to the bouquets of the links. In that setting, one can take Heegaard triples which are either left or right subordinate, corresponding to putting the top degree generator on the right or the left side in the Heegaard triple. In the case of band maps, there is no ambiguity of whether a band is an α-band or a β-band, and hence there is no ambiguity in whether a triple is right or left subordinate to a band, as there was in triples subordinate to a link. As an analog of switching between right and left subordinate diagrams for a framed link, we will prove a relation between the α-band maps and the β-band maps in Proposition 7.8.

We also make the following remark:

Remark 6.9. One can work with Heegaard triples that were obtained from triples subordinate to a band by winding some of the curves to achieve admissibility. We note that using the basic properties of the band maps, as well as the fact that the top degree generators $\Theta^{\mathbf{w}}_{\alpha\beta}$ or $\Theta^{\mathbf{w}}_{\beta\gamma}$ are defined as elements of homology (which doesn't depend on the choice of diagrams), one can use such triples for which we have wound some of the $\boldsymbol{\alpha}, \boldsymbol{\beta}$ or $\boldsymbol{\gamma}$ curves in order to achieve admissibility to compute band maps.

6.3. **Maps for type-w bands.** Suppose that B^α is an α-band for a link $\mathbb{L}$ in Y. Let w and w' be the $\mathbf{w}$-basepoints adjacent to the ends of B^α. Again by Lemma 3.2, there is a distinguished generator as a $\mathbb{Z}_2[U_{\mathbf{w}}]$-module

$$\Theta^{\mathbf{z}}_{\alpha\beta} \in \mathcal{H}FL^-_{\max,\mathbf{z}}(\mathcal{H}_{\alpha\beta}, \sigma, \mathfrak{P}, \mathfrak{s}_0) \subseteq \mathcal{H}FL^-(\mathcal{H}_{\alpha\beta}, \sigma, \mathfrak{P}, \mathfrak{s}_0)$$

where $\mathcal{H}FL^-_{\max,\mathbf{z}}$ denotes the subgroup of top homological degree with respect to the grading $\mathrm{gr}_{\mathbf{z}}$, as long as $(\sigma, \mathfrak{P})$ assigns the same color to w and w'.

Suppose that $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple subordinate to a band. For $\mathfrak{s} \in \mathrm{Spin}^c(Y)$ we can define

$$F^{\mathbf{w}}_{B^\alpha}(\cdot) = F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}(\Theta^{\mathbf{z}}_{\alpha\beta} \otimes \cdot) : \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}).$$

Analogously for a β-band B^β, we can pick a triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ subordinate to B^β and define

$$F^{\mathbf{w}}_{B^\beta}(\cdot) = F_{\boldsymbol{\alpha\beta\gamma},\mathfrak{s}}(\cdot \otimes \Theta^{\mathbf{z}}_{\beta\gamma}) : \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z}, \mathfrak{s}) \to \mathcal{C}FL^\circ(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z}, \mathfrak{s}),$$

the map which counts holomorphic triangles with $\mathfrak{s}_{\mathbf{w}}(\psi) = \mathfrak{s} \in \mathrm{Spin}^c(X_{\alpha\beta\gamma})_0 \cong \mathrm{Spin}^c(Y)$. Again $\mathrm{Spin}^c(X_{\alpha\beta\gamma})_0$ denotes the Spin^c structures which restrict to the torsion one on $Y_{\beta\gamma}$.

The same argument used for the type-$\mathbf{z}$ band maps can be used to show that the type-$\mathbf{w}$ yield well defined maps up to equivariant, filtered chain homotopy on $\mathcal{C}FL^\circ$.

7. Properties and relations involving the band maps

In this section, we prove some important properties of the band maps.

7.1. **Band maps and quasi-stabilization.** We now consider trying to commute the band maps and the quasi-stabilization maps. One can often do this, but there are several cases where the bands don't commute, as one would expect, given the interpretation of band maps in terms of dividing sets on surfaces.

Proposition 7.1. *Suppose that B is a band for the link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and w and z are two new basepoints which are adjacent in $L \setminus (\mathbf{w} \cup \mathbf{z})$. Then w follows z, then*

$$F^{\mathbf{z}}_B S^{\pm}_{w,z} \simeq S^{\pm}_{w,z} F^{\mathbf{z}}_B, \qquad \text{if } w \text{ follows } z, \text{ and}$$
$$F^{\mathbf{z}}_B S^{\pm}_{z,w} \simeq S^{\pm}_{z,w} F^{\mathbf{z}}_B, \qquad \text{if } z \text{ follows } w.$$

If the pair (w, z) does not separate the ends of the band B from either of $\mathbf{z}$-basepoints adjacent to B, then

$$F^{\mathbf{z}}_B T^{\pm}_{w,z} \simeq T^{\pm}_{w,z} F^{\mathbf{z}}_B, \qquad \text{if } w \text{ follows } z, \text{ and}$$
$$F^{\mathbf{z}}_B T^{\pm}_{z,w} \simeq T^{\pm}_{z,w} F^{\mathbf{z}}_B, \qquad \text{if } z \text{ follows } w.$$

Proof. This will follow from the triangle map computations of performed in Theorem 5.3, Proposition 5.6, and Lemma 5.1 and (this last one is [MO10, Proposition 5.2]). Let us first consider the case that $B = B^\alpha$ is an α-band. Pick a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ which is subordinate to B^α.

Let us first consider the case that z follows w. In this case, the adjacency assumption required when considering the $T^{\pm}_{z,w}$ maps is automatically satisfied since the basepoints w and z cannot be in the component of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which contain the ends of B^α (as those components are going from $\mathbf{w}$-basepoints to $\mathbf{z}$-basepoints). The maps $S^{\pm}_{z,w}$ and $T^{\pm}_{z,w}$ are computed using a β-quasi-stabilization, i.e. we add a small α_0 curve, and a β_s curve which is allowed to travel around the diagram. Since the pair (z, w) is not adjacent to the ends of B^α, there is a basepoint $w_0 \in \mathbf{w}$, adjacent to z on $L \setminus B^\alpha$. We can now quasi-stabilize the triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{w}, \mathbf{z})$ by adding a small α_0 and β_0 encircling z and w, and a γ_s which encircles z and w_0, as in the top left diagram of Figure 7.1. Now applying Lemma 5.1, we know that for appropriately stretched almost complex structure, with respect to the decomposition given by the $\mathrm{gr}_{\mathbf{w}}$-grading of the new intersection points on the quasi-stabilized triple, the quasi-stabilized band map takes the form

$$F^{\mathbf{z}}_{B^\alpha} = \begin{pmatrix} F^{\mathbf{z}}_{B^\alpha} & 0 \\ 0 & F^{\mathbf{z}}_{B^\alpha} \end{pmatrix}.$$

Hence $F^{\mathbf{z}}_{B^\alpha}$ commutes with both $S^{\pm}_{z,w}$ and $T^{\pm}_{z,w}$.

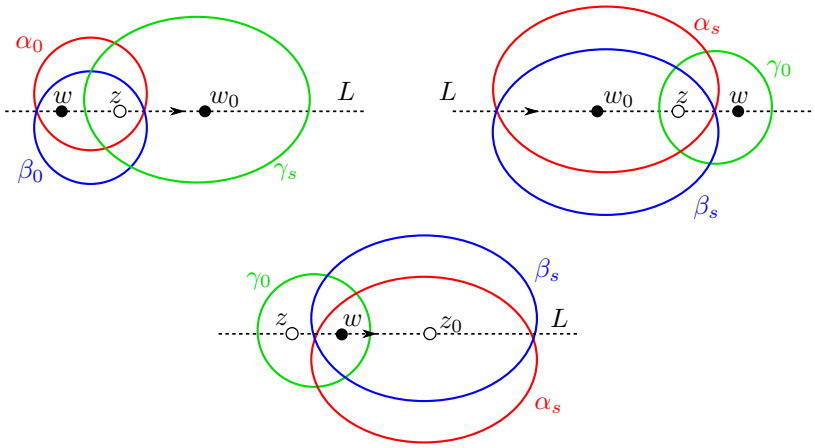

FIGURE 7.1. Three configurations of quasi-stabilizations considered in Proposition 7.1. The dashed lines represent shadows of the link L.

We now consider the case that w follows z. Write $\boldsymbol{\alpha} = \{\alpha_1, \dots, \alpha_n\}$ and $\boldsymbol{\beta} = \{\beta_1, \dots, \beta_n\}$ where $\alpha_1, \dots, \alpha_{n-1}$ are small Hamiltonian isotopies of $\beta_1, \dots, \beta_{n-1}$, and α_n and β_n are as in the construction of the band maps.

Suppose first that the pair (w, z) is not between the ends of the band B^α and one of z_1 and z_2. In this case, there is basepoint $w_0 \in \mathbf{w}$ such that w_0 is adjacent to z on $L \setminus B^\alpha$. In this case, we can form the quasi-stabilized diagram for the band map as in the top right diagram of Figure 7.1. The quasi-stabilized triple shown there is now a triple subordinate to the band B^α, added to the quasi-stabilized link. We claim that $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w})$ is strongly positive with respect to the point p_s^+, the higher $\mathrm{gr}_{\mathbf{w}}$ graded generator of $\alpha_s \cap \beta_s$. To see this, we apply [Zem16, Lemma 6.3], which shows that if $\mathcal{H}$ is a diagram with a chosen point $p^+ \in \mathbb{T}_\alpha \cap \mathbb{T}_\beta$, and $\overline{\mathcal{H}}$ is the diagram obtained by performing surgery on a 0-sphere, then $\mathcal{H}$ is strongly positive at p^+ iff $\overline{\mathcal{H}}$ is strongly positive at p^+. For $i \in \{1, \dots, n-1\}$, the pair of curves α_i and β_i determine an embedded 2-sphere. Upon surgering all of these out and forgetting about the $\mathbf{z}$-basepoints, there are two possibilities for the resulting diagram, depending on the configuration of the basepoints.

The first subcase is that w_0 is adjacent on $(L, \mathbf{w}, \mathbf{z})$ to an end of B^α. Given the ordering of w and z, this is equivalent to the pair (w, z) being adjacent to an end of B^α. The second subcase is that the basepoint w_0 is not adjacent to the ends of B^α, which is equivalent to the pair (w, z) not being adjacent to an end of the band.

In the first subcase, where w_0 is adjacent to an end of B^α, the resulting diagram is

$$(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w_0, w_0', w)$$

shown in Figure 7.2, where w_0' is the other $\mathbf{w}$-basepoint adjacent to the end of B^α. The curves α_n and β_n are isotopic relative the $\mathbf{w}$-basepoints in the diagram, and hence determine a framed 2-sphere. Upon surgering this out, we are left with the diagram $(S^2, \alpha_s, \beta_s, w, w_0)$, which is clearly strongly positive at p_s^+.

The other subcase we need to consider is if w_0 is not adjacent to an end of B^α. Given the ordering of the basepoints w and z, this is equivalent to (w, z) not being adjacent to the ends of the band B^α. then the resulting diagram is

$$(S^2, \alpha_n, \beta_n, w_1, w_2) \sqcup (S^2, \alpha_s, \beta_s, w_0, w)$$

where w_1 and w_2 are now the two basepoints adjacent to the ends of B^α. The curves α_n and β_n each intersect each other twice, and similarly α_s and β_s both intersect each other twice.

Consider now the case shown in Figure 7.2, where (w, z) is between the basepoint $w_0 \in \mathbf{w}$ and an end of the band B^α. In the second case, where w_0 is not adjacent to an end of B^α, the resulting diagram is the disjoint union of two diagrams on genus zero Heegaard surfaces, each with a pair of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ curves intersecting exactly twice. This diagram is obviously strongly positive at p_s^+. Hence the original diagram

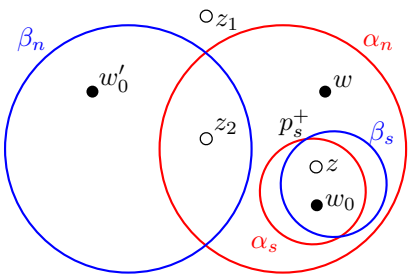

FIGURE 7.2. This is a picture of the first subcase, where (w, z) is between an end of B^α and the basepoint w_0. Once we surger out the 2-spheres determined by the pairs α_i and β_i for $1 \leq i \leq n-1$, we are left with the diagram $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$ shown. The basepoints z_1 and z_2 are the $\mathbf{z}$-basepoints adjacent to the ends of B^α.

was as well, since it can be obtained by surgering on a collection of 0-spheres. In both cases, the triangle computation of Theorem 5.3 implies that the quasi-stabilized triangle map for $F^{\mathbf{z}}_{B^\alpha}$ will take the form

$$F^{\mathbf{z}}_{B^\alpha} = \begin{pmatrix} F^{\mathbf{z}}_{B^\alpha} & 0 \\ 0 & F^{\mathbf{z}}_{B^\alpha} \end{pmatrix},$$

and hence will commute with both the $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ maps.

We now have to consider the final case, which is when the pair (w, z) is between an end of the band B^α and one of the $\mathbf{z}$-basepoints adjacent to the ends of B^α. Call this basepoint z_0. In this case, we can quasi-stabilize near z_0, as in the bottom diagram of in Figure 7.1. If p_s^+ is the higher $\mathrm{gr}_{\mathbf{w}}$-graded intersection point of $\alpha_s \cap \beta_s$, we note that the diagram $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\})$ is no longer strongly positive at p_s^+. Indeed we proceed as before. The proof of [Zem16, Lemma 6.3] (giving a procedure for constructing strongly positive diagrams) applies equally well to weak positivity, instead of just strong positivity, and hence we can surger out all of the 2-spheres determined by the pairs α_i and β_i for $1 \leq i \leq n-1$, leaving the diagram $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$, for basepoints $w_1, w_2 \in \mathbf{w}$. The diagram $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$ is shown in Figure 7.3, and once we forget about the $\mathbf{z}$-basepoints, we can view it as a free-stabilization at w with the curves α_s and β_s.

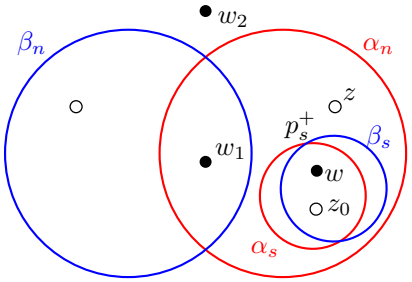

FIGURE 7.3. In the case that (w, z) is between an end of B^α and a $\mathbf{z}$ basepoint, once we surger out the 2-spheres determined by the pairs α_i and β_i for $1 \leq i \leq n-1$, we are left with the diagram $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$.

By Lemma 5.5, the pair $(S^2, \{\alpha_n, \alpha_s\}, \{\beta_n, \beta_s\}, w, w_1, w_2)$ is thus weakly positive at p_s^+. As we remarked, surgering out the 2-spheres induced by the pairs α_i and β_i doesn't affect weak positivity, by the same argument as [Zem16, Lemma 6.3], and hence we conclude that the original diagram $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_s\}, \boldsymbol{\beta} \cup \{\beta_s\}, \mathbf{w} \cup \{w\})$ is weakly positive at p_s^+. By Proposition 5.6, we can thus write

$$F^{\mathbf{z}}_{B^\alpha} = \begin{pmatrix} F^{\mathbf{z}}_{B^\alpha} & * \\ 0 & F^{\mathbf{z}}_{B^\alpha} \end{pmatrix},$$

with respect to the decomposition given by higher and lower generators, with respect to the grading $\mathrm{gr}_{\mathbf{w}}$. This now commutes with $S^{\pm}_{w,z}$ since the bottom left entry is zero, but not necessarily $T^{\pm}_{w,z}$ since the top right entry is not zero.

If B is instead a β-band, the above proof adapts using the obvious adaptations of the triangle lemmas involved obtained by switching the roles of the $\boldsymbol{\alpha}$-, $\boldsymbol{\beta}$- and $\boldsymbol{\gamma}$-curves appropriately. □

Remark 7.2. Notice that the case which was not covered by the previous proposition is when the pair (w,z) separates an end of B from one of the $\mathbf{z}$-basepoints adjacent to B. An interesting exercise is to check that given the interpretation of these maps in terms of surfaces with divides, one would not expect these maps to commute.

By symmetry (or just rewriting the above proof), we have the following:

Proposition 7.3. *Suppose that B is a band for the link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$, and w and z are two new basepoints which are adjacent in $L \setminus (\mathbf{w} \cup \mathbf{z})$. Then w follows z, then*

$$F^{\mathbf{w}}_B S^{\pm}_{w,z} \simeq S^{\pm}_{w,z} F^{\mathbf{w}}_B, \qquad \text{if } w \text{ follows } z, \text{ and}$$

$$F^{\mathbf{w}}_B S^{\pm}_{z,w} \simeq S^{\pm}_{z,w} F^{\mathbf{w}}_B, \qquad \text{if } z \text{ follows } w.$$

If the pair (w,z) does not separate the ends of the band B from either of $\mathbf{w}$-basepoints adjacent to B, then

$$F^{\mathbf{w}}_B T^{\pm}_{w,z} \simeq T^{\pm}_{w,z} F^{\mathbf{w}}_B, \qquad \text{if } w \text{ follows } z, \text{ and}$$

$$F^{\mathbf{w}}_B T^{\pm}_{z,w} \simeq T^{\pm}_{z,w} F^{\mathbf{w}}_B, \qquad \text{if } z \text{ follows } w.$$

7.2. **Band maps and disk stabilization.** The band maps correspond to index 1 critical points on the surface, and the maps $S^{\pm}_{\mathbb{U},D}$ correspond to index 0 and index 2 critical points. In this section we prove that these maps can cancel each other, in an appropriate sense.

Suppose that $\mathbb{U} = (U, w, z)$ is an unknot with two basepoints, spanned by the disk $D \subseteq Y$, which is disjoint from L. Suppose that we attach a band B to $L \cup U$. The surface $B \cup D$, yields an isotopically unique diffeomorphism

$$\phi : (Y, L) \to (Y, (L \cup U)(B)),$$

which is the identity outside of a neighborhood of $B \cup D$. Let $\overline{B}$ denote the dual band, attached to $(Y, L \cup U)(B)$, yielding the link $L \cup U$. We prove the following:

Proposition 7.4. *Let B and $\overline{B}$ be the bands described in the previous paragraph. If B is an α-band then we have*

$$F^{\mathbf{z}}_B S^{+}_{(U,w,z),D} \simeq S^{+}_{w,z}\phi_* \qquad \text{and} \qquad F^{\mathbf{w}}_B S^{+}_{(U,w,z),D} \simeq T^{+}_{w,z}\phi_*.$$

Similarly,

$$S^{-}_{(U,w,z),D} F^{\mathbf{z}}_{\overline{B}} \simeq S^{-}_{w,z}\phi_* \qquad \text{and} \qquad S^{-}_{(U,w,z),D} F^{\mathbf{z}}_{\overline{B}} \simeq T^{-}_{w,z}\phi_*.$$

If B is a β-band, we have

$$F^{\mathbf{z}}_B S^{+}_{(U,w,z),D} \simeq S^{+}_{z,w}\phi_* \qquad \text{and} \qquad F^{\mathbf{w}}_B S^{+}_{(U,w,z),D} \simeq T^{+}_{z,w}\phi_*,$$

and

$$S^{-}_{(U,w,z),D} F^{\mathbf{w}}_{\overline{B}} \simeq T^{-}_{z,w}\phi_* \qquad \text{and} \qquad S^{-}_{(U,w,z),D} F^{\mathbf{w}}_{\overline{B}} \simeq T^{-}_{z,w}\phi_*.$$

The core of the argument are two triangle map computations. Consider the Heegaard triple shown in Figure 7.4, which we view as a modified free-stabilization of a Heegaard triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ at a point $z' \in \mathbf{z}$. In the following two triangle map lemmas, we don't require that $\boldsymbol{\alpha}'$ be related to $\boldsymbol{\alpha}$ in any way (other than assuming the triple assumes the appropriate admissibility assumption). Since we will only apply the lemma in the case that $\boldsymbol{\alpha}'$ is a result of a small Hamiltonian isotopy of $\boldsymbol{\alpha}$, we label the first set of attaching curves as $\boldsymbol{\alpha}'$ for expository purposes. Additionally, recall that for the $\mathbf{z}$-band maps to be defined, we need the $\mathbf{z}$-basepoints adjacent to the end of the band to have the same color. As such, in the following two Lemmas, we will identify both V_z and $V_{z'}$ with a formal variable V.

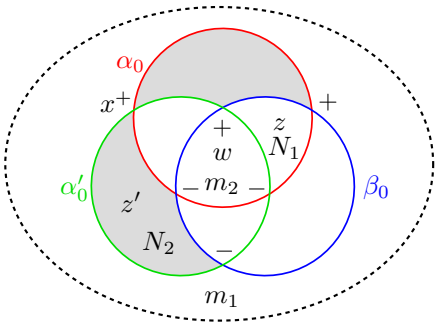

FIGURE 7.4. A modified free-stabilization of a Heegaard triple $(\Sigma, , \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ at a point $z' \in \mathbf{z}$. Multiplicities and integers m_1, m_2, N_1 and N_2 are shown. This is considered in Lemma 7.5. Triangles like this can be used to compute the α-band map for a band which merges an unknot with two basepoints onto a link component. The two shaded regions are each examples of small triangles, which could be counted.

Lemma 7.5. *If $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with $z' \in \mathbf{z}$, and $\overline{\mathcal{T}}$ denotes the stabilization shown in Figure 7.4, near the point z', then if $\mathfrak{s} \in \mathrm{Spin}^c(X_{\alpha'\alpha\beta})$ we have identifications*

$$F_{\overline{\mathcal{T}},\overline{\mathfrak{s}}}(\mathbf{x} \times x^+, \mathbf{y} \times \cdot) = \begin{pmatrix} F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) & * \\ 0 & V \cdot F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) \end{pmatrix},$$

for almost complex structure sufficiently stretched along the dashed circle, and we have identified $V = V_z = V_{z'}$. Here $x^+ \in \alpha_0' \cap \alpha_0$, where here the matrix is an expansion in the intersection point $\cdot \in \alpha_0 \cap \beta_0$, with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading.

Proof. The proof is similar to the proof of [Zem15b, Theorem 6.17]. For a homology triangle $\psi_1 \# \psi_2 \in \pi_2(\mathbf{x} \times x^+, \mathbf{y} \times y, \mathbf{z} \times z)$, we have that

$$\mu(\psi_1 \# \psi_2) = \mu(\psi_1) + 2m_2 + \Delta_{\mathbf{w}}(y, z),$$

where $\Delta_{\mathbf{w}}(y, z)$ denotes the drop in $\mathrm{gr}_{\mathbf{w}}$ grading from y to z. By stretching the neck sufficiently we can ensure that $\mu(\psi_1)$ and $2m_2$ are nonnegative. In particular if $\Delta_{\mathbf{w}}(y, z) = 1$, there are no holomorphic triangles representing a $\psi_1 \# \psi_2$ with Maslov index 0, so the lower left entry of the matrix is zero. We now consider the diagonal entries, where $\Delta_{\mathbf{w}}(y, z) = 0$. We are left with $\mu(\psi_1) = m_2 = 0$. Using the identical argument as in the proof of [Zem15b, Theorem 6.17], we see that for fixed y and z in the same grading, the number of points in $\mathcal{M}(\psi_1)$ is equal to the total number of points in moduli spaces of the form $\mathcal{M}(\psi_1 \# \psi_2)$ where ψ_2 ranges over triangles in $\pi_2(x^+, y, z)$ with $m_2(\psi_2) = 0$. We must now just pay attention to the multiplicities of these triangles at the $\mathbf{z}$-basepoints. Using the multiplicities in Figure 7.4, for a homology triangle $\psi_2 \in \pi_2(x^+, y^+, z^+)$, we claim that

$$m_1 + m_2 = N_1 + N_2.$$

To see this, note that this equation holds for small triangle in $\pi_2(x^+, y^+, z^+)$ shown in Figure 7.4 (the top small triangle). Any other homology element in $\pi_2(x^+, y^+, z^+)$ in this diagram can be obtained by splicing in α_0'-, α_0- or β_0-boundary degenerations of Maslov index $+2$ or -2 (the obvious disks with boundary on α_0', α_0 or β_0, with multiplicity $+1$ or -1, respectively), which changes both $m_1 + m_2$ and $N_1 + N_2$ both by 1 or both by -1, respectively. From this observation the equality $m_1 + m_2 = N_1 + N_2$ follows for any triangle in $\pi_2(x^+, y^+, z^+)$. The triangle ψ_1 has $V_{z'}$ multiplicity m_1, and we thus conclude that any $\psi_1 \# \psi_2$ which is counted has total V_z and $V_{z'}$ multiplicity $N_1 + N_2 = m_1 + m_2 = m_1$, from which we conclude the top left entry of the matrix is as stated.

Analogously, to see that the bottom right entry is as stated, we observe that if $\psi_2 \in \pi_2(x^+, y^-, z^-)$, then

$$N_1 + N_2 = m_1 + m_2 + 1,$$

since it holds for the formula holds for the stated small triangle, and respects splicing in boundary degeneration homology classes. Using the same reasoning as before, we conclude that the bottom right entry has an extra factor of $V = V_z = V_{z'}$. □

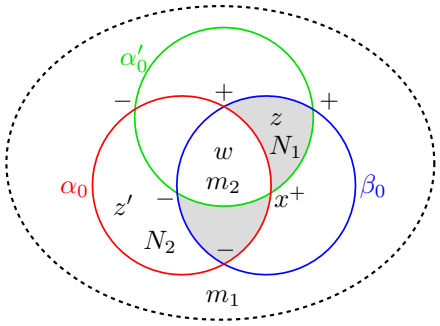

FIGURE 7.5. Another modified free-stabilization of a Heegaard triple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ at a point $z' \in \mathbf{z}$. Multiplicities m_1, m_2, N_1 and N_2 are shown. This is considered in Lemma 7.6, and can be used to compute an α-band map for an α-band which splits off an unknot with two basepoints from a link component. The two shaded regions are each small triangles which might be counted.

Analogously, we will also have use for the following lemma:

Lemma 7.6. *Suppose that $\mathcal{T} = (\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a Heegaard triple with $z' \in \mathbf{z}$, and $\overline{\mathcal{T}}$ denotes the stabilization shown in Figure 7.5, near the point z'. Then if $\mathfrak{s} \in \mathrm{Spin}^c(X_{\alpha'\alpha\beta})$, we have an identification*

$$F_{\overline{\mathcal{T}},\overline{\mathfrak{s}}}(\mathbf{x} \times x^+, \mathbf{y} \times \cdot) = \begin{pmatrix} V \cdot F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) & * \\ 0 & F_{\mathcal{T},\mathfrak{s}}(\mathbf{x}, \mathbf{y}) \end{pmatrix},$$

for almost complex structure sufficiently stretched along the dashed circle, assuming we have identified $V = V_z = V_{z'}$. Here $x^+ \in \alpha_0' \cap \alpha_0$ and the matrix expansion is in the intersection point $\cdot \in \alpha_0 \cap \beta_0$, with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading.

Proof. The proof is nearly identical to the proof of Lemma 7.5. When we stretch the almost complex structure, the same degenerations occur as in that Lemma, but the only change is the multiplicity of these triangles over the $\mathbf{z}$-basepoints. If $\psi_2 \in \pi_2(x^+, y^+, z^+)$, then we have

$$N_1 + N_2 = m_1 + m_2 + 1,$$

by the same argument as before (one checks that the formula holds for a small triangle, then check that it respects splicing in Maslov index ± 2 boundary degenerations). Using the same reasoning as in Lemma 7.5, we conclude that the entry in the top left corner has an extra factor of V_z. Similarly, if $\psi_2 \in \pi_2(x^+, y^-, z^-)$, then we have

$$N_1 + N_2 = m_1 + m_2,$$

and so the same reasoning as before shows that the each holomorphic triangle counted in the bottom right entry will have the same combined V_z and $V_{z'}$ multiplicity as the corresponding triangle had in $F_{\mathcal{T},\mathfrak{s}}$. □

Proof of Proposition 7.4. We first prove that $F_B^{\mathbf{z}} S^+_{(U,w,z),D} \simeq S^+_{w,z}\phi_*$ in the case that B is an α-band. This follows from the model triangle map computation in Lemma 7.5, as follows.

Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ be a Heegaard diagram for $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$. Let $\boldsymbol{\alpha}'$ be small Hamiltonian isotopies of the $\boldsymbol{\alpha}$ curves. We can perform disk stabilization to get a diagram $(\Sigma, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\}, \mathbf{w} \cup \{w\}, \mathbf{z} \cup \{z\})$. We can assume that the band B is embedded in Σ, and define α_0' to be as in Figure 7.4. The triple $(\Sigma, \boldsymbol{\alpha}' \cup \{\alpha_0'\}, \boldsymbol{\alpha} \cup \{\alpha_0\}, \boldsymbol{\beta} \cup \{\beta_0\})$ is a triple subordinate to the band B. By Lemma 7.5, we can write

$$F_B^{\mathbf{z}} = \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} & * \\ 0 & V \cdot \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha} \to \boldsymbol{\alpha}'} \end{pmatrix},$$

where V is identified with both V_z and $V_{z'}$ (recall we need the complexes colored so that z and z' have the same color in order for the band maps to be defined).

Noting that to compute $S^+_{w,z}$, we need to stretch the almost complex structure around an α_0 curve so that the diagram is quasi-stabilized at w and z. Using the change of almost complex structure computation from Lemma 3.23, we know that we can ensure that this map takes the form

$$\Phi_{J\to J'} = \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix}.$$

Hence we have that

$$\Phi_{J\to J'} F^{\mathbf{z}}_B S^+_{(U,w,z),D} = \begin{pmatrix} 1 & * \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} & * \\ 0 & V\cdot \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} \end{pmatrix} \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} \Phi^{\boldsymbol{\beta}}_{\boldsymbol{\alpha}\to\boldsymbol{\alpha}'} \\ 0 \end{pmatrix},$$

which is tautologically just $S^+_{w,z}\phi_*$.

To show $S^-_{(U,w,z),D}F^{\mathbf{z}}_B \simeq S^-_{w,z}\phi_*$, one follows similar reasoning, though one needs to use the triangle map computation from Lemma 7.6. The statement about $\mathbf{w}$-bands follows by symmetry, and the statements about β-bands follow from an easy modification. □

7.3. **Commuting of band maps I.** In this subsection, we prove that band maps can always be commuted, if one is an α-band, and the other is a β-band. We will later consider commuting two band maps when both bands are α-bands or both are β-bands.

Proposition 7.7. *Suppose that B^α_1 is an α-band and B^β_2 is a β-band such that $\mathbb{L}(B^\alpha_1), \mathbb{L}(B^\beta_2)$ and $\mathbb{L}(B^\alpha_1, B^\beta_2)$ all have at least two basepoints on each component. Then, for any choice of $\mathbf{o}_1, \mathbf{o}_2 \in \{\mathbf{w}, \mathbf{z}\}$, we have*

$$F^{\mathbf{o}_1}_{B^\alpha_1} F^{\mathbf{o}_2}_{B^\beta_2} \simeq F^{\mathbf{o}_2}_{B^\beta_2} F^{\mathbf{o}_1}_{B^\alpha_1}.$$

Proof. This follows from the associativity relations for Heegaard Floer homology. Pick curves as follows. Let $(\Sigma, \alpha_1, \dots, \alpha_{n-1}, \beta_1, \dots, \beta_{n-1}, \mathbf{w}, \mathbf{z})$ represent $Y \setminus N(L \cup B^\alpha_1 \cup B^\beta_2)$. Let $\alpha'_1, \dots, \alpha'_{n-1}$ be small Hamiltonian isotopies of the $\alpha_1, \dots, \alpha_{n-1}$ curves, and let $\beta'_1, \dots, \beta'_{n-1}$ be small Hamiltonian isotopies of the $\beta_1, \dots, \beta_{n-1}$ curves. Pick remaining curves $\alpha_n, \beta_n, \alpha'_n$ and β'_n such that

- $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L})$
- $(\Sigma, \alpha_1, \dots, \alpha_n, \beta'_1, \dots, \beta'_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\beta_2))$
- $(\Sigma, \alpha'_1, \dots, \alpha'_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\alpha_1))$,
- and $(\Sigma, \alpha'_1, \dots, \alpha'_n, \beta'_1, \dots, \beta'_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\alpha_1, B^\beta_2))$.

Let us write $\boldsymbol{\alpha}' = (\alpha'_1, \dots, \alpha'_n)$ and $\boldsymbol{\beta}' = (\beta'_1, \dots, \beta'_n)$. Let $\Theta^{\mathbf{o}_1}_{\alpha'\alpha}$ denote the algebraic generator of $\mathcal{H}FL^-_{\max,\mathbf{o}_1}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ and let $\Theta^{\mathbf{o}_2}_{\beta\beta'}$ be the algebraic generator of $\mathcal{H}FL^-_{\max,\mathbf{o}_2}(\Sigma, \boldsymbol{\beta}, \boldsymbol{\beta}')$.

Noting that

$$\mathrm{Spin}^c(X_{\alpha'\beta\beta'})_0 \cong \mathrm{Spin}^c(X_{\alpha'\alpha\beta})_0 \cong \mathrm{Spin}^c(X_{\alpha'\alpha\beta'})_0 \cong \mathrm{Spin}^c(X_{\alpha\beta\beta'})_0 \cong \mathrm{Spin}^c(Y),$$

where $\mathrm{Spin}^c(X)_0$ denotes the Spin^c structures which restrict to the torsion one on the appropriate boundary component corresponding to $(S^1\times S^2)^{\#k}$, we have by the associativity relations for the Heegaard Floer triangle maps that

$$F_{\alpha'\beta\beta',\mathfrak{s}}(F_{\alpha'\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_1}_{\alpha'\alpha}, \mathbf{x}), \Theta^{\mathbf{o}_2}_{\beta\beta'}) \simeq F_{\alpha'\alpha\beta',\mathfrak{s}}(\Theta^{\mathbf{o}_1}_{\alpha'\alpha}, F_{\alpha\beta\beta',\mathfrak{s}}(\mathbf{x}, \Theta^{\mathbf{o}_2}_{\beta\beta'}))$$

for a choice of $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, thus showing that

$$F^{\mathbf{o}_1}_{B^\alpha_1} F^{\mathbf{o}_2}_{B^\beta_2} \simeq F^{\mathbf{o}_2}_{B^\beta_2} F^{\mathbf{o}_1}_{B^\alpha_1}.$$

□

7.4. **Band maps and basepoint moving maps.** We now consider the relation between α-bands and β-bands. Whether a band B is an α-band or a β-band is determined by the configuration of the ends of B with respect to the basepoints, and hence a band B can never be both an α-band and a β-band. On the other hand, given an α-band B^α, we can isotope the ends of B^α across the adjacent $\mathbf{z}$-basepoints to get a β-band B^β. Let z_1 and z_2 denote the adjacent $\mathbf{z}$-basepoints. There is a well defined diffeomorphism $\phi : (Y, \mathbb{L}(B^\alpha)) \to (Y, \mathbb{L}(B^\beta))$ which fixes all basepoints in $\mathbf{w} \cup \mathbf{z}$ except z_1 and z_2. The diffeomorphism ϕ switches z_1 and z_2. This is illustrated in Figure 7.6.

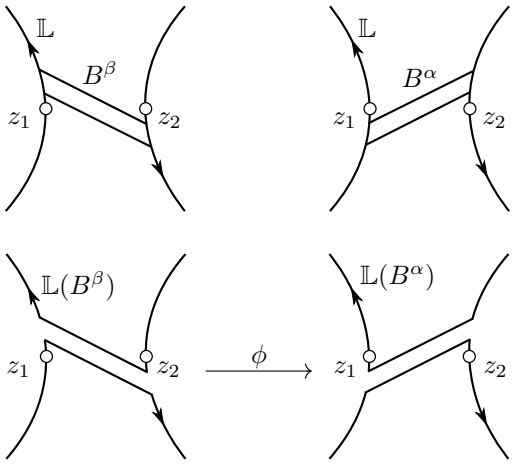

FIGURE 7.6. Two bands B^α and B^β related by an isotopy across $\mathbf{z}$-basepoints, and the diffeomorphism ϕ between $(Y, \mathbb{L}(B^\alpha))$ and $(Y, \mathbb{L}(B^\beta))$, which maps z_1 to z_2 and maps z_2 to z_1.

Proposition 7.8. *If B^α, B^β and ϕ are as in the previous paragraph, then*

$$\phi_* F^{\mathbf{z}}_{B^\alpha} \simeq F^{\mathbf{z}}_{B^\beta}.$$

As a step towards proving the above proposition, consider the following situation. Suppose that D is a disk embedded in $Y \setminus L$, w and z are basepoints on $\partial D = U$, and B^α_0, B^β_0 are bands connecting D to $\mathbb{L}$, as in Figure 7.7. Define diffeomorphisms ϕ_1 and ϕ_2, as in Figure 7.7. Also let ϕ_0 be the basepoint moving map, shown in Figure 7.7.

Note that to prove Proposition 7.8, using diffeomorphism invariance of the band maps, it is sufficient to show

$$(\phi_0)_* F^{\mathbf{z}}_{B^\alpha_0} (\phi_2)_* \simeq F^{\mathbf{z}}_{B^\beta_0} (\phi_1)_*. \tag{5}$$

To this end, we first prove the following:

Lemma 7.9. *The diffeomorphisms ϕ_0, ϕ_1 and ϕ_2 induce maps*

$$(\phi_0)_* \simeq T^-_{z_1,w} S^+_{w,z_2}, \qquad (\phi_1)_* \simeq T^-_{z_1,w} F^{\mathbf{z}}_{B^\alpha_0} S^+_{(U,w,z),D}, \qquad \text{and} \qquad (\phi_2)_* \simeq T^-_{w,z_2} F^{\mathbf{z}}_{B^\beta_0} S^+_{(U,w,z),D}.$$

Proof. The formula for $(\phi_0)_*$ follows immediately Lemma 3.27.

We now consider the formula for $(\phi_1)_*$. Let $\phi^\alpha : (Y, L) \to (Y, (L \cup U)(B^\alpha_0))$ be a diffeomorphism which is fixed outside of a neighborhood of $D \cup B^\alpha_0$. By Proposition 7.4 we have that

$$F^{\mathbf{z}}_{B^\alpha_0} S^+_{(U,w,z),D} \simeq S^+_{w,z} \phi^\alpha_*.$$

Post-composing with $T^-_{z_1,w}$ we get

$$T^-_{z_1,w} F^{\mathbf{z}}_{B^\alpha_0} S^+_{(U,w,z),D} \simeq T^-_{z_1,w} S^+_{w,z} \phi^\alpha_*.$$

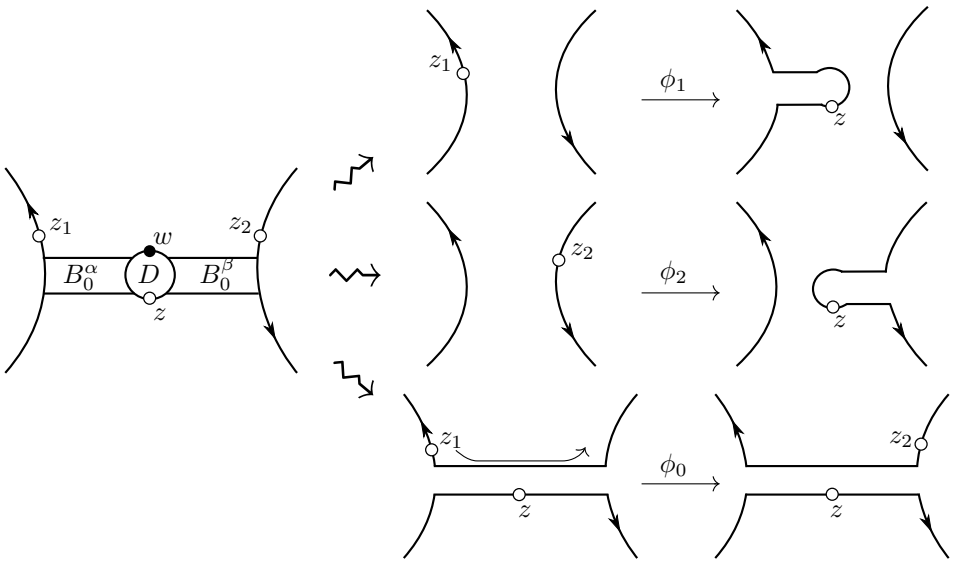

FIGURE 7.7. The disk D, the bands B_0^α and B_0^β, and the diffeomorphisms ϕ_0, ϕ_1 and ϕ_2.

By Lemma 3.27, the composition $T^-_{z_1,w}S^+_{w,z}$ is equivalent to the basepoint moving map $(\phi_{z\to z_1})_*$. Since $\phi_{z\to z_1}\circ\phi^\alpha$ is clearly the diffeomorphism ϕ_1, we thus have

$$T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}S^+_{(U,w,z),D} \simeq (\phi_{z\to z_1})_*\phi^\alpha_* \simeq (\phi_1)_*.$$

The claim about $(\phi_2)_*$ follows from identical reasoning to $(\phi_1)_*$. □

We can now prove Proposition 7.8:

Proof of Proposition 7.8. As we remarked earlier, it is sufficient to show Equation (5). Using Lemma 7.9, we thus need to show

$$T^-_{z_1,w}S^+_{w,z_2}F^{\mathbf{z}}_{B_0^\alpha}T^-_{w,z_2}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D} \simeq F^{\mathbf{z}}_{B_0^\beta}T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}S^+_{(U,w,z),D}.$$

We note that

$$\begin{aligned} T^-_{z_1,w}S^+_{w,z_2}F^{\mathbf{z}}_{B_0^\alpha}T^-_{w,z_2}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D} &\simeq T^-_{z_1,w}S^+_{w,z_2}T^-_{w,z_2}F^{\mathbf{z}}_{B_0^\alpha}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D} && \text{(Lemma 7.1)}\\ &\simeq T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D} && \text{(Lemma 7.7).}\end{aligned}$$

On the other hand we have

$$\begin{aligned} F^{\mathbf{z}}_{B_0^\beta}T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}S^+_{(U,w,z),D} &\simeq T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\beta}F^{\mathbf{z}}_{B_0^\alpha}S^+_{(U,w,z),D} && \text{(Lemma 7.1)}\\ &\simeq T^-_{z_1,w}F^{\mathbf{z}}_{B_0^\alpha}F^{\mathbf{z}}_{B_0^\beta}S^+_{(U,w,z),D} && \text{(Proposition 7.7).}\end{aligned}$$

These two expressions are clearly equal, completing the proof. □

Remark 7.10. We note that the proof of Proposition 7.8 can be interpreted similarly to [OS06, Lemma 5.2], which says that one can can compute maps induced surgery on framed links using either left subordinate or right subordinate Heegaard triples. The argument uses the fact that the 3-handle maps cancel both the left subordinate, and right subordinate maps, for 0-surgery on an unlink. Analogously, our proof of Proposition 7.8 follows from a similar observation that the addition of an unknot via the maps $S^+_{\mathbb{U},D}$ can be canceled by either the α-band maps, or the β-band maps.

The analogous statement also holds for the $\mathbf{w}$ band maps. Suppose that B^α is an α-band and B^β is a β-band resulting by sliding the ends of B^α across the two adjacent $\mathbf{w}$-basepoints (call them w_1 and w_2). Then there is a well defined diffeomorphism $\phi : (Y, \mathbb{L}(B^\alpha)) \to (Y, \mathbb{L}(B^\beta))$ which moves w_1 to w_2 along the

link (as in Figure 7.6, but with basepoints w_1 and w_2 replacing the z_1 and z_2 which appear in that picture.) We have the following:

Proposition 7.11. *If B^α, B^β and ϕ are as in the previous paragraph, then*

$$F^{\mathbf{w}}_{B^\alpha} = \phi_* F^{\mathbf{w}}_{B^\beta}.$$

The proof is, of course, to just rewrite the proof of Proposition 7.8, switching Ss with Ts and $\mathbf{w}$s with $\mathbf{z}$s.

7.5. **Commuting band maps II.** We now revisit the question of commuting two bands maps in the case that both are α-bands, or both are β-bands. We phrase the result for α-bands, but of course the proof carries over without change to β-bands.

Proposition 7.12. *Suppose that B_1 and B_2 are two α-bands on the link $\mathbb{L} = (L, \mathbf{w}, \mathbf{z})$ such that $\mathbb{L}(B_1), \mathbb{L}(B_2)$ and $\mathbb{L}(B_1, B_2)$ all have at least two basepoints on each component. Then we have*

$$F^{\mathbf{z}}_{B_1} F^{\mathbf{z}}_{B_2} \simeq F^{\mathbf{z}}_{B_2} F^{\mathbf{z}}_{B_1},$$

and

$$F^{\mathbf{w}}_{B_1} F^{\mathbf{w}}_{B_2} \simeq F^{\mathbf{w}}_{B_2} F^{\mathbf{w}}_{B_1}.$$

Suppose further that there exist points $\mathbf{d} \subseteq L \setminus (\mathbf{w} \cup \mathbf{z})$ such that each adjacent pair of points in $\mathbf{w} \cup \mathbf{z}$ is separated by exactly one point of $\mathbf{d}$. If B_2 has its ends in regions of $L \setminus \mathbf{d}$ containing $\mathbf{w}$-basepoints, and B_1 has its ends in regions of $L \setminus \mathbf{z}$ containing $\mathbf{z}$-basepoints, then

$$F^{\mathbf{z}}_{B_1} F^{\mathbf{w}}_{B_2} \simeq F^{\mathbf{w}}_{B_2} F^{\mathbf{z}}_{B_1}.$$

We will give two proofs of this fact. The first is very fast, and mostly unenlightening. The second is longer, but more enlightening.

Proof 1. Suppose that B_1 and B_2 are two bands, and suppose that we assign type-$\mathbf{z}$ to B_1 and type-$\mathbf{o}_2$ to B_2. The only case of the proposition which isn't covered is the case that we were to assign a designation of type-$\mathbf{w}$ to both B_1 and B_2, though this case is an easy modification of the case where B_1 and B_2 are both assigned type-$\mathbf{z}$.

It is easily verified the assumptions of the Proposition statement imply that there are two $\mathbf{z}$-basepoints which are adjacent to the ends of one of the bands on both $\mathbb{L}$ and $\mathbb{L}(B_2)$. In the case that both B_1 and B_2 are given type-$\mathbf{z}$, we let B_1 denote this band. In the case that only B_1 is of type-$\mathbf{z}$, the configuration of the points in $\mathbf{d}$ forces this band to be B_1. Let B_1' be the β-band obtained by sliding B_1 across the $\mathbf{z}$-basepoints adjacent to B_1. Let $\phi : L(B_1) \to L(B_1')$ be the obvious modification of the diffeomorphism pictured in Figure 7.6. Note that ϕ fixes the band B_2, so we can define a diffeomorphism $\phi^{B_2} : L(B_1)(B_2) \to L(B_1')(B_2)$, which is equal to ϕ as a map on Y. We now apply Propositions 7.7 and 7.8. We have that

$$\begin{aligned} F^{\mathbf{z}}_{B_1} F^{\mathbf{o}_2}_{B_2} &\simeq \phi^{B_2}_* F^{\mathbf{z}}_{B_1'} F^{\mathbf{o}_2}_{B_2} \\ &\simeq \phi^{B_2}_* F^{\mathbf{o}_2}_{B_2} F^{\mathbf{z}}_{B_1'} \\ &\simeq F^{\mathbf{o}_2}_{B_2} \phi_* F^{\mathbf{z}}_{B_1'} \\ &\simeq F^{\mathbf{o}_2}_{B_2} F^{\mathbf{z}}_{B_1}, \end{aligned}$$

completing the proof. Note that $\phi^{B_2}_* F^{\mathbf{o}_2}_{B_2} \simeq F^{\mathbf{o}_2}_{B_2} \phi_*$ since ϕ fixes the band B_2. □

We offer an alternate proof, which is somewhat more enlightening and more in the spirit with other standard proofs of this type from Heegaard Floer homology. The reader can safely ignore this if they are satisfied with the first proof, though we include it to give an alternate perspective.

Proof 2. We now offer a second proof, which is somewhat more explicit. There are three cases to consider:

(1) B_1 and B_2 have ends on four separate components of $L \setminus (\mathbf{w} \cup \mathbf{z})$;
(2) B_1 and B_2 each have exactly one end on the same component of $L \setminus (\mathbf{w} \cup \mathbf{z})$, but the other ends are on other components;
(3) B_1 and B_2 each have ends in a total of two components of $L \cup (\mathbf{w} \cup \mathbf{z})$.

We first show that the case (2) can be reduced to case (1), and case (3) can be reduced to case (2). Consider case (2). On the component of $L \setminus (\mathbf{w} \cup \mathbf{z})$ which contains an end of B_1 and an end of B_2, let $d \in \mathbf{d}$ denote a point separating the bands. We let z and w denote new basepoints near a, ordered so that w comes after z. Let $\mathbf{o}_1, \mathbf{o}_2 \in \{\mathbf{w}, \mathbf{z}\}$ denote the type designations of B_1 and B_2 respectively. If $\mathbf{o}_1 = \mathbf{o}_1 = \mathbf{z}$, then we note that $S^-_{w,z}$ commutes with $F^{\mathbf{z}}_{B_i}$ by Proposition 7.1, and hence we simply compute, using the result for case (1), that

$$\begin{aligned}
F^{\mathbf{z}}_{B_1} F^{\mathbf{z}}_{B_2} &\simeq F^{\mathbf{z}}_{B_1} F^{\mathbf{z}}_{B_2} S^-_{w,z} T^+_{w,z} && \text{(Lemma 3.13)} \\
&\simeq S^-_{w,z} F^{\mathbf{z}}_{B_1} F^{\mathbf{z}}_{B_2} T^+_{w,z} && \text{(Proposition 7.1)} \\
&\simeq S^-_{w,z} F^{\mathbf{z}}_{B_2} F^{\mathbf{z}}_{B_1} T^+_{w,z} && \text{(Case (1))} \\
&\simeq F^{\mathbf{z}}_{B_2} F^{\mathbf{z}}_{B_1} S^-_{w,z} T^+_{w,z} && \text{(Proposition 7.1)} \\
&\simeq F^{\mathbf{z}}_{B_2} F^{\mathbf{z}}_{B_1} && \text{(Lemma 3.13).}
\end{aligned}$$

The same strategy can be used to reduce Case (3) to Case (2). Using $T^-_{w,z}$ instead of $S^-_{w,z}$ establishes the result in the case that both $\mathbf{o}_1 = \mathbf{o}_2 = \mathbf{w}$. In the case that $\mathbf{o}_1 = \mathbf{w}$ and $\mathbf{o}_2 = \mathbf{z}$, we let w and z be new basepoints as before. We note that the pair (w, z) does not separate B_1 from the $\mathbf{w}$ basepoints adjacent to the ends of B_1, as can be seen in Figure 7.8. Applying Propositions 7.1 and 7.3 shows that $S^-_{w,z}$ commutes with both $F^{\mathbf{w}}_{B_1}$ and $F^{\mathbf{z}}_{B_2}$ so we can reduce from Case (2) to Case (1), and from Case (3) to Case (2), as before.

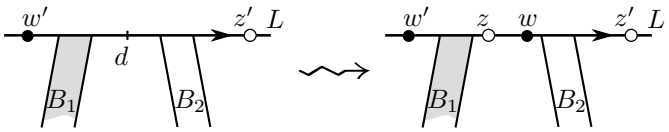

FIGURE 7.8. The ends of the bands B_1 and B_2, and the adjacent $\mathbf{w}$-basepoint w'. Also shown are the new basepoints w and z, and the point d between the ends of B_1 and B_2.

Hence we need only to show the statement in the case that the ends of B_1 and B_2 are in separate components of $L \setminus (\mathbf{w} \cup \mathbf{z})$. Note that in this case, the condition of the Proposition statement about a set of points $\mathbf{d} \subseteq L \setminus (\mathbf{w} \cup \mathbf{z} \cup B_1 \cup B_2)$ is trivially satisfied (it was only necessary to ensure that the quasi-stabilization maps commuted appropriately, so we could reduce to this case). The remainder of the proof now follows from the associativity relations, and a model computation, as we now describe.

We pick a Heegaard diagram $(\Sigma, \alpha_1, \dots, \alpha_{n-2}, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ which induces a sutured diagram for the sutured manifold $Y \setminus N(L \cup B^\alpha_1 \cup B^\alpha_2)$, with sutures induced by the basepoints $\mathbf{w}$ and $\mathbf{z}$. For $i \in \{1, \dots, n-2\}$, let α^1_i, α^2_i and α^{12}_i be small isotopies of α_i, in such a way that each pair of curves with the same index i intersects in exactly two points, but curves with differing indices do not intersect (e.g. $|\alpha_i \cap \alpha^1_j| = 2\delta_{ij}$).

Let α_{n-1} and α_n be curves such that $(\Sigma, \alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $\mathbb{L}$. Now pick curves α^τ_{n-1} and α^τ_n for $\tau \in \{1, 2, 12\}$ such that

- $(\Sigma, \alpha^1_1, \dots, \alpha^1_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\alpha_1))$;
- $(\Sigma, \alpha^2_1, \dots, \alpha^2_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\alpha_2))$;
- $(\Sigma, \alpha^{12}_1, \dots, \alpha^{12}_n, \beta_1, \dots, \beta_n, \mathbf{w}, \mathbf{z})$ represents $(Y, \mathbb{L}(B^\alpha_1, B^\alpha_2)$;
- α^2_n is a small isotopy of α_n;
- α^1_{n-1} is a small isotopy of α_{n-1};
- α^{12}_{n-1} is a small isotopy of α^2_{n-1};
- α^{12}_n is a small isotopy of α^1_n.

Define $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)$, $\boldsymbol{\alpha}^1 = (\alpha^1_1, \dots, \alpha^1_n)$ and $\boldsymbol{\alpha}^2$ and $\boldsymbol{\alpha}^{12}$ similarly. Note that

- $(\Sigma, \boldsymbol{\alpha}^1, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is subordinate the band B^α_1, attached to the link $\mathbb{L}$;
- $(\Sigma, \boldsymbol{\alpha}^{12}, \boldsymbol{\alpha}^1, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is subordinate to the band B^α_2, attached to $\mathbb{L}(B^\alpha_1)$;
- $(\Sigma, \boldsymbol{\alpha}^2, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is subordinate to band B^α_2, attached to $\mathbb{L}$;
- $(\Sigma, \boldsymbol{\alpha}^{12}, \boldsymbol{\alpha}^2, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is subordinate to the band B^α_1, attached to $\mathbb{L}(B^\alpha_2)$.

We note that by construction $(\Sigma, \boldsymbol{\alpha}^1, \boldsymbol{\alpha}, \boldsymbol{\beta})$, $(\Sigma, \boldsymbol{\alpha}^{12}, \boldsymbol{\alpha}^1, \boldsymbol{\beta})$, $(\Sigma, \boldsymbol{\alpha}^2, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and $(\Sigma, \boldsymbol{\alpha}^{12}, \boldsymbol{\alpha}^2, \boldsymbol{\beta})$ all represent unlinks in (disjoint unions of) $(S^1 \times S^2)^{\#k}$ where each component of the unlink has exactly two basepoint, except for one component, which has four. As such, the associated $\mathcal{H}FL^-_{\max,\mathbf{o}_i}$ groups have a generator of top degree. We write

$$\Theta^{\mathbf{o}_1}_{\alpha_1\alpha} \in \mathcal{H}FL^-_{\max,\mathbf{o}_1}(\Sigma, \boldsymbol{\alpha}_1, \boldsymbol{\alpha}),$$
$$\Theta^{\mathbf{o}_2}_{\alpha_2\alpha} \in \mathcal{H}FL^-_{\max,\mathbf{o}_2}(\Sigma, \boldsymbol{\alpha}_2, \boldsymbol{\alpha}),$$
$$\Theta^{\mathbf{o}_2}_{\alpha_{12}\alpha_1} \in \mathcal{H}FL^-_{\max,\mathbf{o}_2}(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}_1),$$
$$\Theta^{\mathbf{o}_1}_{\alpha_{12}\alpha_2} \in \mathcal{H}FL^-_{\max,\mathbf{o}_1}(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}_2),$$

for these algebraic generators.

For $\mathfrak{s} \in \mathrm{Spin}^c(Y)$, the associativity relations for the quadruple $(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}_1, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ read

$$F_{\alpha_{12}\alpha_1\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_2}_{\alpha_{12}\alpha_1}, F_{\alpha_1\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_1}_{\alpha_1\alpha}, \mathbf{x})) \simeq F_{\alpha_{12}\alpha\beta,\mathfrak{s}}(F_{\alpha_{12}\alpha_1\alpha,\mathfrak{s}_0}(\Theta^{\mathbf{o}_2}_{\alpha_{12}\alpha_1}, \Theta^{\mathbf{o}_1}_{\alpha_1\alpha}), \mathbf{x}).$$

Similarly the associativity relations for the quadruple $(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}_2, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ read

$$F_{\alpha_{12}\alpha_2\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_1}_{\alpha_{12}\alpha_2}, F_{\alpha_2\alpha\beta,\mathfrak{s}}(\Theta^{\mathbf{o}_2}_{\alpha_2\alpha}, \mathbf{x})) \simeq F_{\alpha_{12}\alpha\beta,\mathfrak{s}}(F_{\alpha_{12}\alpha_2\alpha,\mathfrak{s}_0}(\Theta^{\mathbf{o}_1}_{\alpha_{12}\alpha_2}, \Theta^{\mathbf{o}_2}_{\alpha_2\alpha}), \mathbf{x}).$$

Hence we have that $F^{\mathbf{o}_1}_{B_1} F^{\mathbf{o}_2}_{B_2} \simeq F^{\mathbf{o}_2}_{B_2} F^{\mathbf{o}_1}_{B_1}$ as long as

$$F_{\alpha_{12}\alpha_1\alpha,\mathfrak{s}_0}(\Theta^{\mathbf{o}_2}_{\alpha_{12}\alpha_1}, \Theta^{\mathbf{o}_1}_{\alpha_1\alpha}) = F_{\alpha_{12}\alpha_2\alpha,\mathfrak{s}_0}(\Theta^{\mathbf{o}_1}_{\alpha_{12}\alpha_2}, \Theta^{\mathbf{o}_2}_{\alpha_2\alpha}),$$

as elements in the homology group $\mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$.

Let $\boldsymbol{\alpha}'$ be small isotopies of the $\boldsymbol{\alpha}$ curves. Let $\widetilde{B}_1$ and $\widetilde{B}_2$ denote the bands B_1 and B_2, now viewed as bands for unlinks inside of $(S^1 \times S^2)^{\#g}$. Let Θ be the top degree generator of $\mathcal{H}FL^-_{\max}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$. Note that $\Theta^{\mathbf{o}_1}_{\alpha_1\alpha}$ is equal, on homology to

$$F^{\mathbf{o}_1}_{\widetilde{B}_1}(\Theta),$$

because tautologically the triangle map which computes that band map can be interpreted, when applied to Θ, as a change of diagrams map applied to the element $\Theta^{\mathbf{o}_1}_{\alpha_1\alpha}$. Analogously, $\Theta^{\mathbf{o}_2}_{\alpha_2\alpha}$ is equal on homology to $F^{\mathbf{o}_2}_{\widetilde{B}_2}(\Theta)$. Upon these considerations, it is thus sufficient to show that

$$F^{\mathbf{o}_1}_{\widetilde{B}_1} F^{\mathbf{o}_2}_{\widetilde{B}_2}(\Theta) = F^{\mathbf{o}_2}_{\widetilde{B}_2} F^{\mathbf{o}_1}_{\widetilde{B}_1}(\Theta),$$

as elements of $\mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}_{12}, \boldsymbol{\alpha})$. This can now be done via a model computation. Note that we just need to verify this on homology, for a single choice of diagrams. By using the triangle map computation of [Zem15b, Theorem 8.8], since we are applying the band maps to Θ, the top degree generator, by surgering out the 2-spheres corresponding to pairs α'_i and α_i for $i \neq n-1, n$, we can assume that Σ is genus zero, and that there are only four components of our link, each with two basepoints, and that the bands connect two pairs of these. Considering the diagram in Figure 7.9, we can eliminate the middle pair of $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ curves by applying [Zem15b, Theorem 8.8] once again. This reduces the problem to the case that the two pairs of unknots are each in different components of $S^3 \sqcup S^3$, and now commutation is trivial.

□

Remark 7.13. We still used some trickery in reducing the complexity of the model computations, in the previous alternate proof, since we used the trick of adding then removing basepoints. If one didn't want to do this, they would need to do model computations for cases (2) and (3) separately. Case (2) would be where there are three unknots, and we want add 2 bands, each with two ends on one of the unknots. One would then need to check that the band maps commuted, when applied to the top degree generator. Case (2) can actually be performed, without too much trouble. One can check the desired commutations on a genus zero diagram, with $|\boldsymbol{\alpha}| = |\boldsymbol{\beta}| = 2$ and $|\mathbf{w}| = |\mathbf{z}| = 3$. Using the triangle counts from [Zem15b, Theorem 8.8] cleverly, one can see that for appropriate almost complex structure, the band maps commute for the configuration described in the above lemma. When the configuration of bands is not as in the previous lemma, the above triangle count actually doesn't apply (we end up counting triangles which correspond to the $*$ component of the triangle map). Case (3) would be to add two bands to two unknots. The resulting link is slightly more complicated. It ends up being a two component link in $S^1 \times S^2$, with each component being isotopic to $S^1 \times \{*\}$, but going in opposite directions. Similar holomorphic triangle counts can be

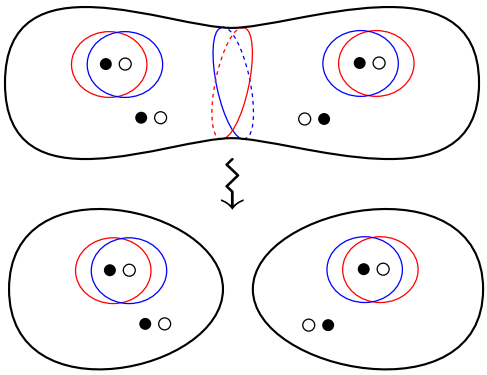

FIGURE 7.9. A diagram $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$ for unknots in S^3 which we can reduce to for our model band computation. By eliminating the middle $\boldsymbol{\alpha}'$ and $\boldsymbol{\alpha}$ curves, we reduce to the case that the bands $\tilde{B}_1$ and $\tilde{B}_2$ are in different components of $S^3 \sqcup S^3$.

used. We've included a picture showing the starting and ending links in both cases (2) and (3), if one were to use this strategy.

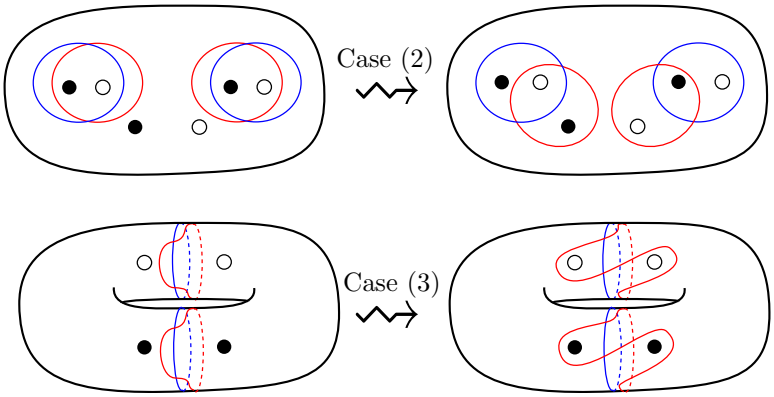

FIGURE 7.10. The remaining model computations needed to prove Proposition 7.12 without using any basepoint tricks. Each of the squiggly arrows corresponds to a composition of two band maps. Each band map moves one of the two $\boldsymbol{\alpha}$ curves. In both cases, for appropriately stretched almost complex structure, the maps can be computed by using [Zem15b, Theorem 8.8], as long as the configuration of the types of the bands is as in the Proposition statement. In the first diagram, we can try to view the diagram as a free stabilization of a diagram with one $\boldsymbol{\alpha}$- and one $\boldsymbol{\beta}$-curve, and in the second diagram, we can try to view the diagram as a 1-handle stabilization of a diagram with with one $\boldsymbol{\alpha}$- and one $\boldsymbol{\beta}$-curve. In each case, there are two ways to view the entire diagram as a stabilization of a smaller diagram. The configuration of basepoints with respect to the ends ensures that one of the configurations can legitimately be viewed as a free stabilization or 1-handle stabilization with respect to one of the gradings $\mathrm{gr}_{\mathbf{w}}$ or $\mathrm{gr}_{\mathbf{z}}$, with two intersection points being the top degree generator. In such a case, we can use the triangle map computation of [Zem15b, Theorem 8.8] to arrive at the expected triangle count. In the case that the bands are both of the "wrong" type for the configuration of ends, when we try to view the diagrams as a free stabilization or 1-handle stabilization, we will be counting triangles with vertices on the bottom generators. One can explicitly find examples where "unexpected" triangles are counted, resulting in the band maps not commuting.

Remark 7.14. We now also describe an alternate route to defining the quasi-stabilization maps. The quasi-stabilization maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ are fairly challenging to define, as they involve challenging holomorphic triangle arguments. An alternate approach to defining them is as follows. One can define the disk stabilization maps $S^{\pm}_{(U,w,z),D}$ and the band maps $F^{\mathbf{o}}_B$. In light of Proposition 7.4, the quasi-stabilization maps are then equal to a composition of disk stabilization, band maps and diffeomorphism maps. Disk stabilization maps commute with each by the same argument that 1-handle maps or free-stabilization maps commute with each other. Commutation (in the appropriate cases) of the quasi-stabilization maps, such as Theorem 3.15, would then follow by using the second proof of Proposition 7.12 (without any of the basepoint tricks, as in Remark 7.13). One would of course need to perform several triangle map computations to show the relation between this alternate definition of the quasi-stabilization maps and the basepoint moving maps. We will not pursue this further, but it's instructive to know that there are other routes to the theory.

7.6. **A final relation between the $F^{\mathbf{z}}_B$ and $F^{\mathbf{w}}_B$ maps.** There is one final relation that the maps $F^{\mathbf{z}}_B$ and $F^{\mathbf{w}}_B$ maps satisfy, assuming a restriction on the coloring.

Proposition 7.15. *Suppose that B is a band (either α- or β-). Suppose that the ends of the band B are adjacent to the basepoints z_1, z_2, w_1 and w_2. If both w_1 and w_2 have the same color, and both z_1 and z_2 have the same color, then*

$$F^{\mathbf{w}}_B \simeq \Psi_{z_1} F^{\mathbf{z}}_B + F^{\mathbf{z}}_B \Psi_{z_1} \simeq \Psi_{z_2} F^{\mathbf{z}}_B + F^{\mathbf{z}}_B \Psi_{z_2}$$

and

$$F^{\mathbf{z}}_B \simeq \Phi_{w_1} F^{\mathbf{w}}_B + F^{\mathbf{w}}_B \Phi_{w_1} \simeq \Phi_{w_2} F^{\mathbf{w}}_B + F^{\mathbf{w}}_B \Phi_{w_2}.$$

Proof. Let us consider the case that B is an α-band; the case that B is a β-band is analogous. We will show the first equivalence. The second follows from symmetry. Let $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ be a triple subordinate to the band B. The diagram $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \mathbf{w}, \mathbf{z})$ is a diagram of an unlink in $(S^1 \times S^2)^{\#g}$, where all but one component of the unlink has exactly two basepoints, and a single component has four. Let $\Theta^{\mathbf{o}}_{\alpha'\alpha} \in \mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ denote the generator of $\mathcal{H}FL^-_{\max,\mathbf{o}}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$. We claim that $\Psi_{z_i}\Theta^{\mathbf{w}}_{\alpha'\alpha} = \Theta^{\mathbf{z}}_{\alpha'\alpha}$ as elements of $\mathcal{H}FL^-(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$, for $i = 1, 2$. As $\Theta^{\mathbf{o}}_{\alpha'\alpha}$ is a well defined element of homology, and the maps Ψ_{z_i} are well defined on the coherent chain homotopy types, it is sufficient to show this relation for a particular choice of diagram and almost complex structure. We prove the statement by induction on the number of attaching curves. Note first that when $\Sigma = S^2$ and $|\boldsymbol{\alpha}'| = |\boldsymbol{\alpha}| = 1$, the claim is obvious, since we simply pick $\boldsymbol{\alpha}'$ and $\boldsymbol{\alpha}$ to intersect at exactly two points, dividing $S^2 \setminus (\boldsymbol{\alpha}' \cup \boldsymbol{\alpha})$ into four bigons, each with a single basepoint. There is a single Maslov index one holomorphic disk which goes over a given z_i, and it has multiplicity one and goes from $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ to $\Theta^{\mathbf{z}}_{\alpha'\alpha}$. More generally, we can proceed by induction to show the claim using the differential computation of [Zem15b, Lemma 8.5]. This implies that if α_n and α'_n are two curves on Σ, small isotopies of each other, and intersecting exactly twice, then the differential can be written as

$$\partial = \begin{pmatrix} \partial_0 & * \\ 0 & \partial_0 \end{pmatrix},$$

for appropriate almost complex structure, where ∂_0 is the differential on the diagram obtained by surgering out α'_n and α_n. The matrix decomposition is in terms of the relative grading of the two intersection points of α'_n and α_n (note that the designation of higher and lower intersection point of $\alpha'_n \cap \alpha_n$ agrees for $\mathrm{gr}_{\mathbf{w}}$ and $\mathrm{gr}_{\mathbf{z}}$ in this case). Hence

$$\Psi_{z_i} = \begin{pmatrix} \Psi_{z_i} & * \\ 0 & \Psi_{z_i} \end{pmatrix}.$$

Noting that in terms of this decomposition, on the complex with α'_n and α_n, the element $\Theta^{\mathbf{o}}_{\alpha'\alpha}$ is equal to $\begin{pmatrix} \Theta^{\mathbf{o}}_{\alpha'\alpha} \\ 0 \end{pmatrix}$, where the entry $\Theta^{\mathbf{o}}_{\alpha'\alpha}$ inside the vector refers to the top $\mathrm{gr}_{\mathbf{o}}$ generator on the complex without the α'_n and α_n curves. Since $\Psi_{z_i}\Theta^{\mathbf{w}}_{\alpha'\alpha} = \Theta^{\mathbf{z}}_{\alpha'\alpha}$ on the complex with α'_n and α_n surgered out, we see that

$$\begin{pmatrix} \Psi_{z_i} & * \\ 0 & \Psi_{z_i} \end{pmatrix} \begin{pmatrix} \Theta^{\mathbf{w}}_{\alpha'\alpha} \\ 0 \end{pmatrix} = \begin{pmatrix} \Theta^{\mathbf{z}}_{\alpha'\alpha} \\ 0 \end{pmatrix},$$

showing that the same relation holds on the complex with α'_n and α_n present.

Gromov compactness yields that

$$F_{\alpha'\alpha\beta} \circ (\partial_{\alpha'\alpha} \otimes 1 + 1 \otimes \partial_{\alpha\beta}) + \partial_{\alpha'\beta} \circ F_{\alpha'\alpha\beta} = 0.$$

Differentiating with respect to the variable V_{z_i}, for either $i=1$ or $i=2$, we see that

$$F_{\alpha'\alpha\beta}\circ(\Psi_{z_i}\otimes 1+1\otimes\Psi_{z_i})+\Psi_{z_i}\circ F_{\alpha'\alpha\beta}+\left(\frac{d}{dV_{z_i}}F_{\alpha'\alpha\beta}\right)\circ(\partial_{\alpha'\alpha}\otimes 1+1\otimes\partial_{\alpha\beta})+\partial_{\alpha'\beta}\circ\left(\frac{d}{dV_{z_i}}F_{\alpha'\alpha\beta}\right)=0.$$

Here $\frac{d}{dV_{z_i}}F_{\alpha'\alpha\beta}$ can be thought of as the map which counts holomorphic triangles, but with the additional factor of $n_{z_i}(\psi)$ for each triangle, and an additional factor of $V_{z_i}^{-1}$, on the entire map. Now color so that $V_{z_1}=V_{z_2}$ and $U_{w_1}=U_{w_2}$ (so that $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ and $\Theta^{\mathbf{w}}_{\alpha'\alpha}$ are both well defined elements of homology), and apply the previous map to elements $\Theta^{\mathbf{w}}_{\alpha'\alpha}\otimes\mathbf{x}$. Noting that $\partial_{\alpha'\alpha}\Theta^{\mathbf{w}}_{\alpha'\alpha}=0$, after coloring, we see that

$$F^{\mathbf{w}}_B=\left(\mathbf{x}\mapsto F_{\alpha'\alpha\beta}(\Theta^{\mathbf{z}}_{\alpha'\alpha},\mathbf{x})\right)=\left(\mathbf{x}\mapsto F_{\alpha'\alpha\beta}(\Psi_{z_i}\Theta^{\mathbf{w}}_{\alpha'\alpha},\mathbf{x})\right)\simeq\Psi_{z_i}F^{\mathbf{z}}_B+F^{\mathbf{z}}_B\Psi_{z_i}.$$

The equivalence $F^{\mathbf{z}}_B\simeq\Phi_{w_i}F^{\mathbf{w}}_B+F^{\mathbf{w}}_B\Phi_{w_i}$ follows by symmetry. □

8. Maps for framed 0-spheres along L

In this section, we define maps for performing surgery on a framed 0-sphere $\mathbb{S}_0$ which occurs along the link L. The maps are defined as a composition involving 1-handle maps away from L, band maps, and maps induced by diffeomorphisms.

The reader may notice that this construction may force us to consider links which are not nullhomologous (in the sense that each component may not be nullhomologous). We give a brief discussion, later in this section.

We begin with the definition:

Definition 8.1. *Suppose L is an oriented link in Y. We say that a framed 0-sphere $\mathbb{S}_0$ in Y, with embedding $\iota: S^0\times D^3\to Y$, is a **framed 0-sphere along** L if the following holds:*

1. *Viewing D^3 as $\{(y,w,z): y^2+w^2+z^2\leq 1\}$, we have that $\iota^{-1}(L)=S^0\times\{(y,0,0): -1\leq y\leq 1\}$;*
2. *Each of $\iota(S^0\times\{(1,0,0\})$ and $\iota(S^0\times\{(-1,0,0)\})$ contains an incoming point of L, and an outgoing point of L, according to the orientation of L.*

8.1. **Model compound 1-handle/band cobordisms.** We introduce a model cobordism for a link cobordism obtained by adding a 4-dimensional 1-handle to $Y\times[0,1]$ with a 2-dimensional 1-handle added to $L\times[0,1]$ inside of the 4-dimensional handle.

Suppose we are given a framed 0-sphere $\mathbb{S}_0$ along $\mathbb{L}$ (with embedding $\iota: S^0\times D^3\to Y$). We form the surgered 3-manifold

$$Y(\mathbb{S}_0)=(Y\setminus\operatorname{int}\iota(S^0\times D^3))\sqcup_\iota([-1,1]\times S^2),$$

using the same construction as if $\mathbb{S}_0$ were a framed 0-sphere embedded away from L.

Inside of $Y(\mathbb{S}_0)$, we can form the link $L(\mathbb{S}_0)$, which is

$$L(\mathbb{S}_0)=(L\setminus\iota(\mathbb{S}_0))\cup_\iota([-1,1]\times\{\pm1,0,0\}).$$

Condition (1) implies that the link $L(\mathbb{S}_0)$ is well defined, while condition (2) implies that $L(\mathbb{S}_0)$ can be given an orientation compatible with L.

Let $W(Y,\mathbb{S}_0)$ denote the 1-handle cobordism defined as

$$W(Y,\mathbb{S}_0)=(Y\times[0,1])\cup(D^1\times D^3)\cup(Y(\mathbb{S})\times[1,2])$$

We define a surface, $\Sigma(\mathbb{S}_0)$, inside of $W(Y,\mathbb{S}_0)$, by

$$\Sigma(\mathbb{S}_0)=(L\times[0,1])\cup B\cup(L(\mathbb{S}_0)\times[1,2])$$

where $B\subseteq[-1,1]\times D^3$ is the subset $B=[-1,1]\times\{(y,w,z): -1\leq y\leq 1, z=w=0\}$.

Note that B is canonically identified with $[-1,1]\times[-1,1]$, and there is a distinguished point $(0,0)$ in B.

8.2. **Construction of the compound 1-handle/band maps.** Suppose that $\mathbb{S}_0$ is a framed 0-sphere along L, which doesn't intersect any of the basepoints in $\mathbf{w} \cup \mathbf{z}$. If $\mathfrak{s} \in \mathrm{Spin}^c(Y)$ and $\overline{\mathfrak{s}}$ denotes the unique extension to the 4-manifold $W(Y, \mathbb{S}_0)$ and $\mathfrak{s}'$ denotes its restriction to $Y(\mathbb{S}_0)$, we will define a type-$\mathbf{z}$ compound 1-handle/band map

$$F^{\mathbf{z}}_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}}} : \mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y(\mathbb{S}_0), \mathbb{L}(\mathbb{S}_0), \sigma', \mathfrak{P}, \mathfrak{s}')$$

as well as a type-$\mathbf{w}$ compound handle/band map

$$F^{\mathbf{w}}_{Y,\mathbb{L},\mathbb{S}_0,\overline{\mathfrak{s}}} : \mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s}) \to \mathcal{C}FL^\circ(Y(\mathbb{S}_0), \mathbb{L}(\mathbb{S}_0), \sigma', \mathfrak{P}, \mathfrak{s}').$$

Our definition of the compound 1-handle/band map is motivated by the following observation. If we pick an element $v \in S^2 \setminus \{(\pm 1, 0, 0)\}$, we can push the surface $\Sigma(\mathbb{S}_0) \subseteq W(Y, \mathbb{S}_0)$ in the direction v towards $[-1, 1] \times S^2 \subseteq [-1, 1] \times D^3$, while keeping the band fixed along $\{(\pm 1, 0, 0)\} \times S^2$, where $[-1, 1] \times D^3$ is the 1-handle of the cobordism, so that the band B becomes a subset of the 3-manifold $Y(\mathbb{S}_0) \times \{1\}$. Thinking of the link cobordism in terms of a sequence of moves of the 4-manifold and surface, this corresponds to first pushing L off of the framed spheres $\mathbb{S}_0$ in the direction of v, then attaching a 4-dimensional 1-handle, and then attaching a band to L (or rather L pushed away from the framed sphere) inside of $Y(\mathbb{S}_0)$, which now goes across the subset $[-1, 1] \times S^2 \subseteq Y(\mathbb{S}_0)$. Such a presentation of the link cobordism is determined by the choice of $v \in S^2 \setminus \{\pm 1, 0, 0\}$. Similarly, a path between any two choices of such v yields an isotopy of the resulting presentation of the link cobordism through such presentations. As $S^2 \setminus \{(\pm 1, 0, 0)\}$ is connected, we conclude that any two presentations are isotopic in this sense. With this observation in mind, we now proceed to formally define the maps.

For the type $\mathbf{o}$ maps (where $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$), we require that the two $\mathbf{o}$-basepoints adjacent to the points in $\mathbb{S}_0$ have the same color.

Given a framed 0-sphere $\mathbb{S}_0$ along L, which doesn't intersect any of the basepoints, we first pick a vector $\theta \in S^1 = \{(0, w, z) : w^2 + z^2 = 1\}$. Such a vector determines a half disk H_i in each $\{p_i\} \times D^3$ (the set of points with spherical coordinates which have angle θ), and an arc a on S^2 (the boundary of the half disk). There is naturally a square $S = [-1, 1] \times a$ determined in $[-1, 1] \times S^2$.

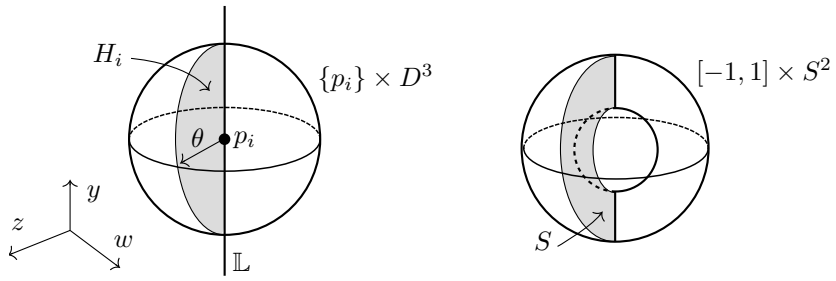

FIGURE 8.1. On the left is the vector $\theta \in S^1$, as well as a half disk H_i, the link $\mathbb{L}$ and the balls $\{p_i\} \times D^3$. On the right is $[-1, 1] \times S^2$, and the square S with boundary on $\{-1, 1\} \times S^2 \cup [-1, 1] \times \{(\pm 1, 0, 0)\}$.

Extend the half disks to slightly larger half disks $\overline{H}_1$ and $\overline{H}_2$ which extend just past the balls $\{p_i\} \times D^3$. The half disks $\overline{H}_1$ and $\overline{H}_2$ determine a diffeomorphism, ψ^θ of Y which pushes $\mathbb{L}$ across the half disks $\overline{H}_i$ to a link $\mathbb{L}_\theta$ just outside of the balls $\{p_i\} \times D^3$ which doesn't intersect $\mathbb{S}_0$, and is the identity outside a neighborhood of the $\{p_i\} \times D^3$. The extensions $\overline{H}_i$ of the disks H_i depend on a contractible set of choices.

Inside of $Y(\mathbb{S}_0)$, there is a natural band

$$B_\theta = \overline{H}_1 \cup S \cup \overline{H}_2$$

from $\mathbb{L}_\theta$ to $\mathbb{L}(\mathbb{S}_0)$.

For $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$. We now define the type-$\mathbf{o}$ compound 1-handle/band to be the composition

$$F^{\mathbf{o}}_{Y,\mathbb{S}_0,\mathbb{L},\overline{\mathfrak{s}}} = F^{\mathbf{o}}_{B_\theta} \circ F_{Y,\mathbb{L}_\theta,\mathbb{S}_0,\overline{\mathfrak{s}}} \circ \psi^\theta_*.$$

Lemma 8.2. *The maps $F^{\mathbf{o}}_{Y,\mathbb{L},\mathbb{S}_0,\bar{\mathfrak{s}}}$ (for $\mathbf{o} \in \{\mathbf{w},\mathbf{z}\}$) are independent of the choice of $\theta \in S^1$, as well as the extensions $\overline{H}_i$.*

Proof. Fix a choice of $\mathbf{o} \in \{\mathbf{w},\mathbf{z}\}$, for the argument. The extended half disks $\overline{H}_i$ depend on a contractible set of choices (a vector field pointing out of the boundary of $\mathbb{S}_0$, and a Riemannian metric), and hence it is easy to see that the maps don't depend on the extensions $\overline{H}_i$.

We now consider dependence on $\theta \in S^1$. Let θ and θ' be two elements in S^1, and let ψ^θ and $\psi^{\theta'}$ be the two associated diffeomorphisms, and let B_θ and $B_{\theta'}$ be the two bands. We pick diffeomorphism Ψ of Y which rotates $\mathbb{L}_\theta$ to $\mathbb{L}_{\theta'}$, which is supported in a neighborhood of $\mathbb{S}_0$, but is constant on $\mathbb{S}_0$. Note that since Ψ is constant on $\mathbb{S}_0$, we can consider the diffeomorphism $\Psi^{\mathbb{S}_0}$ on $Y(\mathbb{S}_0)$.

Note that $\Psi^{\mathbb{S}_0}$ sends the band B_θ to a band which is isotopic relative its boundary to $B_{\theta'}$. As a consequence, we conclude that

$$F^{\mathbf{o}}_{B_{\theta'}} \circ \Psi^{\mathbb{S}_0}_* \simeq \Psi^{\mathbb{S}_0}_* \circ F^{\mathbf{z}}_{B_\theta}\,, \tag{6}$$

by diffeomorphism invariance of the band maps.

Furthermore, since Ψ fixes $\mathbb{S}_0$, we have

$$\Psi^{\mathbb{S}_0}_* \circ F_{Y,\mathbb{L}_\theta,\mathbb{S}_0,\bar{\mathfrak{s}}} \simeq F_{Y,\mathbb{L}_{\theta'},\mathbb{S}_0,\bar{\mathfrak{s}}} \circ \Psi_*. \tag{7}$$

Note too that

$$\Psi \circ \psi^\theta \sim \psi^{\theta'}, \tag{8}$$

where $\sim$ denotes isotopy relative $\mathbb{L}$.

Finally, we note that $\Psi^{\mathbb{S}_0}$ is isotopic to the identity relative $\mathbb{L}(\mathbb{S}_0)$, so

$$\Psi^{\mathbb{S}_0}_* \simeq 1 \tag{9}$$

on $\mathcal{C}FL^\circ(Y(\mathbb{S}_0),\mathbb{L}(\mathbb{S}_0),\mathfrak{s}')$. Hence

$$\begin{aligned}
F^{\mathbf{o}}_{B_\theta} F_{Y,\mathbb{L}_\theta,\mathbb{S}_0,\bar{\mathfrak{s}}}\psi^\theta_* &\simeq (\Psi^{\mathbb{S}_0}_*)^{-1} F^{\mathbf{o}}_{B_{\theta'}} \Psi^{\mathbb{S}_0}_* F_{Y,\mathbb{L}_\theta,\mathbb{S}_0,\bar{\mathfrak{s}}}\psi^\theta_* && \text{(Equation (6))}\\
&\simeq (\Psi^{\mathbb{S}_0}_*)^{-1} F^{\mathbf{o}}_{B_{\theta'}} F_{Y,\mathbb{L}_{\theta'},\mathbb{S}_0,\bar{\mathfrak{s}}}\Psi_*\psi^\theta_* && \text{(Equation (7))}\\
&\simeq (\Psi^{\mathbb{S}_0}_*)^{-1} F^{\mathbf{o}}_{B_{\theta'}} F_{Y,\mathbb{L}_{\theta'},\mathbb{S}_0,\bar{\mathfrak{s}}}\psi^{\theta'}_* && \text{(Equation (8))}\\
&\simeq F^{\mathbf{o}}_{B_{\theta'}} F_{Y,\mathbb{L}_{\theta'},\mathbb{S}_0,\bar{\mathfrak{s}}}\psi^{\theta'}_*, && \text{(Equation (9))}
\end{aligned}$$

completing the proof. □

8.3. **A remark on non-nullhomologous links and Alexander gradings.** We remark that surgering on a framed zero sphere on a link L may produce link components which are not nullhomologous, though the sum of the link components does remain nullhomologous. It's not immediately obvious how to modify the construction of link Floer homology from [OS08] to handle this.

In the original theory, if Y is an integer homology sphere, then $\widehat{CFL}(Y,\mathbb{L})$ has an Alexander multi-grading over $\mathbb{Z}^\ell$, where ℓ is the number of components of a link. For links in manifolds where not all of the components are nullhomologous, such a multi-grading cannot be defined.

We note that in general, there is no problem defining the modules $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ even if the links are not nullhomologous. When the Alexander grading is defined on $\widehat{CFL}(Y,\mathbb{L},\mathfrak{s})$, one can extend this to all of $\mathcal{C}FL^\infty(Y,\mathbb{L},\mathfrak{s})$ by requiring U_w to be -1 graded in component of the grading corresponding to the link component containing w, and defining the variables V_z to be $+1$ graded in the corresponding component.

For non-nullhomologous links, a full $\mathbb{Z}^\ell$ grading cannot be defined, but a $\mathbb{Z}^k$ grading can sometimes be defined for a $k < n$ by collapsing components of the grading. This will be addressed in greater detail in [Zema]. Suppose that $\pi : L \to P$ is a map, which we call a *grading assignment*. We will say that L is *P-nullhomologous* if $\pi^{-1}(p) = [0] \in H_1(Y;\mathbb{Z})$ for each $p \in P$. To a P-nullhomologous knot, we can define a relative Alexander multi-grading on $\mathcal{C}FL^\circ(Y,\mathbb{L},\mathfrak{s})$ over $\mathbb{Z}^P$. This can be upgraded to an absolute grading when $\mathfrak{s}$ is torsion.

A natural question to ask is when the link cobordism maps are graded with respect to the Alexander multi-grading. Naturally our answer involves colorings of links and surfaces. Suppose that $\pi : C(\Sigma) \to P$ is

a map, where $C(\Sigma)$ denotes the components of Σ. The map π naturally induces a map on the ends of Σ. A coloring $(\sigma, \mathfrak{P})$ is P-gradable if

$$\sigma(s) = \sigma(t) \qquad \Rightarrow \qquad \pi(s) = \pi(t)$$

and

$$\sigma(\Sigma_{\mathbf{w}}) \cap \sigma(\Sigma_{\mathbf{z}}) = \varnothing.$$

The above definition naturally adapts to give a definition of a P-gradable coloring of a link $\mathbb{L}$, which is compatible with a grading assignment $\pi : L \to P$. If $(\sigma, \mathfrak{P})$ is a P-gradable coloring, the colored modules $\mathcal{C}FL^\circ(Y, \mathbb{L}, \sigma, \mathfrak{P}, \mathfrak{s})$ have an Alexander multi-grading over $\mathbb{Z}^P$, and the link cobordism maps satisfy a grading change formula (see [Zema]).

We note that if L is P-nullhomologous, and we attach a band B between two link components which have the same value under π, then $L(B)$ is also P-nullhomologous. Similarly if $(\sigma, \mathfrak{P})$ is a P-gradable coloring, and the ends of a type $\mathbf{o}$ band are adjacent to two type $\mathbf{o}$ basepoints of the same color (here $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$), then the induced coloring on $\mathbb{L}(B)$ is also P-gradable.

We now proceed to perform a model computation, to see how the complex $\mathcal{C}FL^\circ$ changes through this operation. If we pick a diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ for $(Y, \mathbb{L})$ and let w, w', z and z' be the basepoints adjacent to the framed zero-sphere $\mathbb{S}$, which is along L. Starting from $\mathcal{H}$, we can add a 1-handle near z and z', and then move the basepoints inside of the bigons on the 1-handle, to get a diagram $\mathcal{H}(\mathbb{S})$ for $(Y(\mathbb{S}), \mathbb{L}(\mathbb{S}))$. This is shown in Figure 8.2.

FIGURE 8.2. From a diagram for $(Y, \mathbb{L})$ (on the left) we can construct a diagram for $(Y(\mathbb{S}), \mathbb{L}(\mathbb{S}))$ (on the right).

The differential on $\mathcal{H}(\mathbb{S})$ is

$$\partial_{\mathcal{H}(\mathbb{S})} = \begin{pmatrix} \partial_{\mathcal{H}} & U_w + U_w \\ V_z + V_{z'} & \partial_{\mathcal{H}} \end{pmatrix}$$

for an appropriate choice of almost complex structure, using the decomposition of higher and lower generators with respect to the $\mathrm{gr}_{\mathbf{w}}$ grading. The diagonal and lower left components can be computed using [Zem15b, Lemma 8.5], and the top right component can be computed using [OS08, Proposition 6.5]. We will not have use for this formula, so we will not include the details of this computation. If we color so that $V_z = V_{z'}$ and $U_w = U_{w'}$, then we note that $\mathcal{C}FL^\circ(\mathcal{H}(\mathbb{S}), \mathfrak{s}\#\mathfrak{s}_0)$ simply becomes $\mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s}) \oplus \mathcal{C}FL^\circ(\mathcal{H}, \mathfrak{s})$, with one copy in the latter expression being shifted up by 1 from the other in the component of the Alexander multi-grading corresponding to the grading assigned to w, w', z and z'.

9. Parametrized Kirby decompositions of link cobordisms

In this section, we define notions of parametrized cobordisms and elementary parametrized cobordism. We then define a way of dividing a link cobordism into elementary parametrized cobordisms, which we call a parametrized Kirby decomposition, borrowing the terminology of [Juh15a].

9.1. Parametrized cobordisms and elementary parametrized cobordisms.

We now define a parametrized link cobordism:

Definition 9.1. *If $(W, \Sigma) : (Y, L) \to (Y', L')$ is an undecorated link cobordism, a **parametrization** of (W, Σ), is a diffeomorphism of link cobordisms*

$$\Phi : W(Y, L, \mathbb{S}) \to (W, \Sigma),$$

where $\mathbb{S}$ is a pairwise disjoint collection of framed spheres in Y (along L or away from L), and $W(Y, L, \mathbb{S})$ is a model link cobordism obtained by attaching 4-dimensional handles, and 2-dimensional handles, as described in Sections 4.1 and 8.1.

The following fact will be useful:

Remark 9.2. If $\mathbb{S}$ and $\mathbb{S}'$ are disjoint collections of framed spheres in Y, the link cobordisms $W(Y, L, \mathbb{S} \cup \mathbb{S}')$, $W(Y(\mathbb{S}), L, \mathbb{S}') \circ W(Y, L, \mathbb{S})$, and $W(Y(\mathbb{S}'), L, \mathbb{S}) \circ W(Y, L, \mathbb{S}')$ are all diffeomorphic. The diffeomorphism is canonically specified up to isotopy fixing $Y \times \{0\}$, and preserving Σ. To see this, it is sufficient to specify the diffeomorphism from the first cobordism to the second. Notice that we can think of the second as being equal to the first with an extra copy of $Y(\mathbb{S}) \times [0,1]$ inserted. A choice of bicollar neighborhood of $Y(\mathbb{S}) \times \{0\}$ allows us to specify such a diffeomorphism, up to isotopy. Since any two bicollar neighborhoods are isotopic, a diffeomorphism between the two cobordisms is specified uniquely up to isotopy fixing $Y \times \{0\}$ and preserving Σ.

Now suppose $(W, \Sigma) : (Y_1, L_1) \to (Y_2, L_2)$ and $(W', \Sigma') : (Y_2, L_2) \to (Y_3, L_3)$ are two link cobordisms with parametrizations

$$\Phi : W(Y_1, L_1, \mathbb{S}_1) \to (W, \Sigma), \qquad \text{and} \qquad \Phi' : W(Y_2, L_2, \mathbb{S}_2) \to (W', \Sigma').$$

For notational simplicity, assume that no components of $\mathbb{S}_1$ or $\mathbb{S}_2$ are framed 0-spheres along L_1 or L_2. Let $\mathbb{S}_2' \subseteq Y_1(\mathbb{S}_1)$ denote $\Phi|^{-1}_{Y_1(\mathbb{S}_1)}(\mathbb{S}_2)$. We can form the **concatenation** of the two parametrizations

$$\Phi' * \Phi : W(Y_1(\mathbb{S}_1), L_1, \mathbb{S}_2') \circ W(Y_1, L_1, \mathbb{S}_1) \to (W', \Sigma') \circ (W, \Sigma),$$

which is defined as Φ on $W(Y_1, L_1, \mathbb{S}_1)$ and $\Phi' \circ \Phi|_{Y_1(\mathbb{S}_1)}$ on $W(Y_1(\mathbb{S}_1), L_1, \mathbb{S}_2')$. If $\mathbb{S}_2' \subseteq Y_1 \setminus \mathbb{S}_1$, then in light of Remark 9.2, the concatenated diffeomorphisms induce a parametrization

$$\Phi' * \Phi : W(Y_1, L_1, \mathbb{S}_1 \cup \mathbb{S}_2') \to (W', \Sigma') \circ (W, \Sigma)$$

which is well defined up to isotopies (of link cobordisms) fixing $Y_1 \times \{0\}$. The concatenation of the two parametrizations can also be defined in the case that $\mathbb{S}_1$ or $\mathbb{S}_2$ is a framed link along L_1 or L_2, with only the obvious notational changes to the domain of $\Phi' * \Phi$.

Analogous to [Juh15b], we need a notion of elementary parametrized cobordisms, corresponding to link cobordisms which have a Morse function with at most one critical value of $f|_{W \setminus \Sigma}$, $f|_\Sigma$ or $f|_{\mathcal{A}}$:

Definition 9.3. *A tuple $\mathcal{W} = (W, F, \Phi, \mathbb{S})$ is called an **elementary, parametrized link cobordism** if $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a decorated link cobordism with $\partial W = -Y_1 \sqcup Y_2$, $F = (\Sigma, \mathcal{A})$ is a surface with divides, and $\mathbb{S}$ is a (possibly empty) collection of framed spheres (0-,1-, or 2-spheres away from L, or 0-spheres along L) in Y_1 which are pairwise disjoint and such that exactly one of the following holds:*

1. *$\mathbb{S} = \varnothing$ and $\Phi : W(Y_1, L_1, \varnothing) \to (W, \Sigma)$ is a parametrization which is the identity on $Y_1 \times \{0\}$ is defined up to isotopies of link cobordisms which fix $Y_1 \times \{0\}$. Furthermore, each arc of the divides $\mathcal{A}$ goes from the incoming boundary to the outgoing boundary.*
2. *$\mathbb{S} \neq \varnothing$ and $\mathbb{S}$ is a framed 0-sphere, 1-dimensional link or 2-sphere away from L_1 and Φ is a parametrization*
$$\Phi : W(Y_1, L_1, \mathbb{S}) \to (W, \Sigma)$$
which is the identity on $Y_1 \times \{0\}$, and which is defined up to isotopies of link cobordisms which fix $Y_1 \times \{0\}$. Furthermore, each arc of the divides $\mathcal{A}$ goes from the incoming boundary to the outgoing boundary.
3. *$\mathbb{S} = \varnothing$ and Φ is a parametrization*
$$\Phi : W(Y_1, L_1, \varnothing) \to (W, \Sigma),$$
which is the identity on $Y_1 \times \{0\}$, and is defined up to isotopies fixing $Y_1 \times \{0\}$. Furthermore, all divides of $\mathcal{A}$, except for exactly one arc A, go from the incoming boundary to the outgoing boundary. There are four subtypes:
   - *I^+: The arc A has two ends on the outgoing boundary, and above the arc is a type-**w** region.*
   - *I^-: The arc A has two ends on the incoming boundary, and below the arc is a type-**w** region.*
   - *II^+: The arc A has two ends on the outgoing boundary, and above the arc is a type-**z** region.*
   - *II^-: The arc A has two ends on the incoming boundary, and below the arc is a type-**z** region.*

   These dividing sets can be found in Figure 11.1.
4. *$\mathbb{S}$ is a framed 0-sphere, embedded along L_1, and Φ is a parametrization*
$$\Phi : W(Y_1, L_1, \mathbb{S}) \to (W, \Sigma),$$

which is the identity on $Y_1 \times \{0\}$, and such that the divides in $\Phi^{-1}(\mathcal{A}) \subseteq \Sigma(\mathbb{S})$ do not intersect $B \subseteq \Sigma(\mathbb{S})$. The parametrization is well defined up to isotopies which fix $Y_1 \times \{0\}$ and the point $(0,0) \in B$ Furthermore, all of the arcs in $\mathcal{A}$ are arcs from the incoming boundary to the outgoing boundary.

*The type (**w** or **z**) of each region of $\Sigma(\mathbb{S}) \setminus \Phi^{-1}(\mathcal{A})$ is well defined. Hence there are two subtypes, (type **w** and type **z**) corresponding to the type of the region with the band under the parametrization.*

Remark 9.4. In the definition of an elementary parametrized link cobordism, the diffeomorphism Φ in the definition is one of undecorated link cobordisms, and doesn't necessarily respect the divides beyond what is stated for type (4) link cobordisms. Furthermore, the parametrization Φ is only defined up to isotopies which fix $Y_1 \times \{0\}$, and in the case of (4), isotopies fixing $(0,0)$.

We make the following definition, borrowing the terminology from [Juh15a] (see also [Juh15b, Definition 2.8]):

Definition 9.5. *For a link cobordism $(W,F) : (Y,\mathbb{L}) \to (Y',\mathbb{L}')$, we define a **parametrized Kirby decomposition**, $\mathcal{K}$, of the decorated link cobordism (W,F) to be the following collection of data:*

(1) A decomposition of (W,F) as

$$(W,F) = (W_0,F_0) \cup_{Y_1} \cdots \cup_{Y_n} (W_n,F_n)$$

where (W_i,F_i) is a link cobordism with $\partial W_i = -Y_i \sqcup Y_{i+1}$ and $\partial F_i = -L_i \sqcup L_{i+1}$, for 3-manifolds with links (Y_i,L_i) and surface with divides F_i (all embedded in W), such that $Y_0 = Y$ and $Y_{n+1} = Y'$ and $\mathcal{A}$ intersects each component of each L_i nontrivially and transversally;

(2) A choice of basepoints $\mathbf{w}_i$ and $\mathbf{z}_i$ on each L_i, with one basepoint per component of $L_i \setminus \mathcal{A}$, such that the type of each basepoint is compatible with the type induced by the divides $\mathcal{A}$.

(3) An embedded sphere or link, $\mathbb{S}_i \subseteq Y_i$ and a parameterizing diffeomorphism Φ_i, making each $(W_i,F_i,\Phi_i,\mathbb{S}_i)$ an elementary, parametrized link cobordism.

Additionally, we assume that all terms of the decomposition of type (2) occur after all of the terms of type (3) or (4).

Finally we assume that the terms of the decomposition of type (2) appear with nondecreasing index (i.e. all 0-spheres occur before 1-dimensional links, which themselves occur before 2-spheres), and there is at most one term of the decomposition which has a 1-dimensional link.

We make the following definition (compare [Juh15b, Section 1.3]):

Definition 9.6. *Suppose $\mathbb{S} \subseteq Y$ is a framed k sphere, and $\mathbb{S}' \subseteq Y(\mathbb{S})$ is a framed $(k+1)$ sphere which intersects the belt sphere of $\mathbb{S}$, transversely at a single point. We say a diffeomorphism*

$$\Psi : W(Y(\mathbb{S}),L_i,\mathbb{S}') \circ W(Y,L,\mathbb{S}) \to W(Y,L,\varnothing),$$

*is **formed by a handle cancellation** if it is constructed by the following procedure: take two Morse functions and gradient like vector fields on $W(Y(\mathbb{S}),L,\mathbb{S}')$ and $W(Y,L,\mathbb{S})$ which each have exactly one critical point, and induce the parametrizations the cobordisms which are the identity on the boundary. Such Morse functions exist by Lemma 10.25. The stable manifold of the higher critical point is transverse to the unstable manifold of the lower critical point and there is a single flowline between them. Using the argument from* [Juh15b, Section 1.3]*, there is a homotopically unique path of functions f_t and gradient like vector fields v_t, which are constant outside a neighborhood of the flowline between the two critical points, which realizes a birth-death singularity. The pair (f_1,v_1) on $W(Y(\mathbb{S}),L,\mathbb{S}') \circ W(Y,L,\mathbb{S})$ determines a diffeomorphism $\Psi : W(Y,L,\varnothing) \to W(Y(\mathbb{S}),L,\mathbb{S}') \circ W(Y,L,\mathbb{S})$, which is the identity on the incoming boundary, $Y \times \{0\}$, and is well defined up to isotopy relative $Y \times \{0\}$, since (f_1,v_1) is determined up to homotopy (through paths of Morse functions and gradient like vector fields, with no intermediate singularities).*

10. Morse theory and moves between parametrized Kirby decompositions

In this section we describe a set of moves that one can be used to connect two parametrized Kirby decompositions. Our strategy is motivated by the observation that a Morse function f_0 on a submanifold can be used to define a Morse function on a regular neighborhood by the formula $f = f_0 + r^2$ (where r is the

distance function induced by a Riemannian metric on the normal bundle). The function f has an index k critical point for each index k critical point of f_0. In this manner, we can construct a Morse function on all of W, by first constructing one on a regular neighborhood of Σ in this manner, and then extending it to all of W. Performing an appropriately adapted Cerf theory will then allow us to move between two presentations given by performing this procedure.

In Subsection 10.1, we list a set of moves to move between two parametrized Kirby decompositions. In Subsection 10.2, we prove several technical results about spaces of coordinate functions, which we will need in order to do the Morse theory. In Subsection 10.3, we show how to nicely connect any two Morse functions which are equal on Σ. In Subsection 10.4, we show how to change the Morse function on Σ. In Subsection 10.5, we show how a (suitably nice) Morse function gives rise to a parametrized Kirby decomposition. Finally in Subsection 10.6, we put everything together, and show that the set of moves from Subsection 10.1 are sufficient to connect any two parametrized Kirby decompositions.

10.1. **Weak equivalences of parametrized Kirby decompositions.** In this section we describe a set of moves to go between two Kirby decompositions.

We first describe a notion of isotopy between two parametrized Kirby decompositions, which will be one of the moves. Suppose that $\mathcal{K}$ is a parametrized Kirby decomposition of (W, F) and that $\phi : W \to W$ is a diffeomorphism which is the identity on ∂W. Then ϕ induces a decomposition $\phi_*\mathcal{K}$ of (W, F) into parametrized cobordisms, though the decomposition may not be a parametrized Kirby decomposition in the sense of Definition 9.5, as we now describe. If

$$(W, F) = (W_0, \Sigma_0) \cup_{Y_1} \cdots \cup_{Y_n} (W_n, \Sigma_n)$$

is the decomposition of (W, Σ) associated to $\mathcal{K}$, then the decomposition associated to $\phi_*\mathcal{K}$ is

$$(\phi(W_0), \phi(\Sigma_0)) \cup_{\phi(Y_1)} \cdots \cup_{\phi(Y_n)} (\phi(W_n), \phi(\Sigma_n)). \tag{10}$$

The parametrizations, Φ_i, of $\mathcal{K}$ can also be pushed forward along ϕ to give diffeomorphisms

$$\phi_*(\Phi_i) : W(\phi(Y_i), \phi(L_i), \phi(\mathbb{S}_i)) \to (\phi(W_i), \phi(\Sigma_i)).$$

The parametrized pieces may fail to be elementary parametrized cobordisms in the sense of Definition 9.3, since the divides $\mathcal{A} \cap \phi_*\Sigma_i$ may not make each $(\phi_*W_i, \phi_*\Sigma_i, \phi_*\Sigma_i \cap \mathcal{A})$ into an elementary parametrized cobordism. With this in mind, we make the following definition:

Definition 10.1. *Suppose that (W, F) is a link cobordism with $F = (\Sigma, \mathcal{A})$. If $\mathcal{K}$ and $\mathcal{K}'$ are two parametrized Kirby decompositions of (W, F), we say that $\mathcal{K}$ and $\mathcal{K}'$ are **strongly isotopic** if there is a 1-parameter family of diffeomorphisms $\phi_t : (W, \Sigma) \to (W, \Sigma)$ such that following holds:*

1. *$\phi_t|_{\partial W} = \mathrm{id}$ for all t;*
2. *$(\phi_t)_*\mathcal{K}$ is a valid parametrized Kirby decomposition for all t;*
3. *$(\phi_1)_*\mathcal{K} = \mathcal{K}'$.*

In the previous definition, note that if ϕ_t is such an isotopy of W and Φ_i is a parameterizing diffeomorphism from a term of $\mathcal{K}$ corresponding to a framed 0-sphere along L_i, then $(\phi_t)_*\Phi_i$ maps the point $(0, 0)$ in the model band to a path on Σ. For ϕ_t to induce a strong isotopy, this path cannot intersect an arc of $\mathcal{A}$, since at the point of time when a point crosses the band, the map $(\phi_t)_*\Phi_i$ would fail to be a valid parametrization for an elementary parametrized cobordism, as it would map the center of the band onto an arc of $\mathcal{A}$.

We now define a more general definition of equivalence between two parametrized Kirby decompositions.

Definition 10.2. *We say two parametrized Kirby decompositions are **weakly equivalent** if they are related by the following set of moves (or their inverses):*

1. *(strong isotopies): See Definition 10.1.*
2. *(isotopy of an attaching sphere, $\mathbb{S}_i$, away from L_i): Suppose $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathbb{S}_i^t$ is a 1-parameter family of framed k-spheres in $Y_i \setminus L_i$ starting at $\mathbb{S}_i$ and ending at a framed sphere $\mathbb{S}_i'$. Let d_i^t denote an extension of this isotopy to Y, we define*

$$D : W(Y_i, L_i, \mathbb{S}_i) \to W(Y_i, L_i, \mathbb{S}_i)$$

by

$$D(x, t) = (d_i^t(x), t)$$

for $(x,t) \in Y \times [0,1]$. We define D to map one handle to the other, under their identifications as $D^{k+1} \times D^{3-k}$. There is an induced diffeomorphism $d_i^{\mathbb{S}_i}$ from $Y_i(\mathbb{S}_i)$ to $Y_i(\mathbb{S}_i')$, which D restricts to. The elementary parametrized cobordism $\mathcal{W}_i$ is replaced with an elementary cobordism $\mathcal{W}_i' = (W_i, F_i, \Phi_i', \mathbb{S}_i')$ where $\Phi_i' \circ D = \Phi_i$.

(3) *(isotopy of a framed 0-sphere sphere $\mathbb{S}_i$ along L_i): Suppose $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathbb{S}_i^t$ is a 1-parameter family of framed 0-spheres along L_i from $\mathbb{S}_i$ to $\mathbb{S}_i'$. The diffeomorphism sending $\mathbb{S}_i$ to $\mathbb{S}_i^t$ extends to an isotopy d_i^t of Y_i which preserves L_i, and sends $\mathbb{S}_i$ to $\mathbb{S}_i^t$. The isotopy d_i^t defines a diffeomorphism*

$$D : W(Y_i, L_i, \mathbb{S}_i) \to W(Y_i, L_i, \mathbb{S}_i')$$

as follows. It is defined by

$$D(x,t) = (d_i^t(x), t),$$

for $(x,t) \in Y \times [0,1]$,

$$D(x,t) = (d_i^1(x), t),$$

for $(x,t) \in Y(\mathbb{S}) \times [1,2]$, and D is the identity on the 1-handles of the cobordisms, under the identification of both as $D^1 \times D^3$. Note that since d_i^t preserves L_i, the map D sends $\Sigma(\mathbb{S})$ to $\Sigma(\mathbb{S}')$, diffeomorphically, and maps the bands to each other as the identity, under the identification of both as $[-1,1] \times [-1,1]$.

The parametrized cobordism $(W_i, F_i, \Phi_i, \mathbb{S}_i)$ is replaced by $\mathcal{W}_i' = (W_i, F_i, \Phi_i', \mathbb{S}_i')$, where $\Phi_i' \circ D = \Phi_i$.

(4) *(adding a trivial cylinder): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ are adjacent terms of $\mathcal{K}$, with at least one of $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ being of type (1), then the pair $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ is replaced with a single $\mathcal{W}_i' = (W_{i+1} \cup W_i, F_{i+1} \cup F_i, \Phi_i', \mathbb{S}_i')$. Under the canonical identification of*

$$W(Y_{i+1}, L_{i+1}, \mathbb{S}_{i+1}) \cup_{\Phi_i} W(Y_i, L_i, \mathbb{S}_i) \cong W(Y_i, L_i, \mathbb{S}_i \cup \Phi_i^{-1}(\mathbb{S}_{i+1})),$$

the maps Φ_i' and the concatenation $\Phi_{i+1} * \Phi_i$ are isotopic through diffeomorphisms fixing $Y_i \times \{0\}$ (or $(Y_i \times \{0\}) \cup \{(0,0)\}$).

(5) *(changing the orientation of a framed sphere): Suppose $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ is an elementary parametrized cobordism with $\mathbb{S}_i \neq \varnothing$, a framed 0-sphere, 2-sphere, or 1-dimensional link, away from L_i. We view $\mathbb{S}_i$ as a map from a $S^k \times D^{3-k}$ to Y_i, and give $S^k \times D^{3-k} \subseteq D^{k+1} \times D^{3-k} \subseteq \mathbb{R}^4$ the coordinates (x,y,w,z) from $\mathbb{R}^4$. Then we can replace $\mathbb{S}_i$ with $\mathbb{S}_i'$ formed by precomposing with the map*

$$\sigma(x,y,w,z) = (-x,y,w,-z).$$

If $\mathbb{S}_i$ is a framed 0-sphere along L_i, we can replace $\mathbb{S}_i$ with the map $\mathbb{S}_i'$ induced by precomposing $\mathbb{S}_i$ with the map

$$\sigma(x,y,w,z) = (-x,-y,w,z).$$

The parameterizing diffeomorphism Φ is also changed by precomposition with σ on $D^{k+1} \times D^{3-k}$.

(6) *(An index 1/2 birth-death singularity away from Σ): Suppose $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ are adjacent terms of the decomposition, both of type (2), such that $\mathcal{W}_i$ has framed 0-sphere $\mathbb{S}_i$ away from L_i and parametrization*

$$\Phi_i : W(Y_i, L_i, \mathbb{S}_i) \to (W_i, \Sigma_i),$$

and $\mathcal{W}_{i+1}$ has framed 1-dimensional link $\mathbb{S}_{i+1}$, and parametrization

$$\Phi_{i+1} : W(Y_{i+1}, L_{i+1}, \mathbb{S}_{i+1}) \to (W_{i+1}, \Sigma_{i+1}).$$

Suppose the belt sphere of $\mathbb{S}_i$ in Y_{i+1} (under the parametrization) intersects a single component, $\mathbb{K}$, of $\mathbb{S}_{i+1}$ exactly once.

The move is to replace $\mathcal{W}_{i+1}$ and $\mathcal{W}_i$ with a pair $\mathcal{W}_{i+1}'$ and $\mathcal{W}_i'$, such that the following holds: There is decomposition of $\mathcal{W}_{i+1}$ as $\mathcal{W}_{i+1}' \circ \mathcal{W}_{i+1}''$ as parametrized cobordisms, where

$$\mathcal{W}_{i+1}' = (W_{i+1}', F_{i+1}', \Phi_{i+1}', \mathbb{S}_{i+1}'), \qquad \text{and} \qquad \mathcal{W}_{i+1}'' = (W_{i+1}'', F_{i+1}'', \Phi_{i+1}'', \mathbb{K})$$

and $(\Phi_{i+1}'')^{-1}(\mathbb{S}_{i+1}') = \mathbb{S}_{i+1} \setminus \mathbb{K}$. The pair $\mathcal{W}_{i+1}'' \circ \mathcal{W}_i$ is replaced by an elementary parametrized cobordism $\mathcal{W}_i' = (W_i', F_i', \Phi_i', \varnothing)$ where

$$\Phi_i' \circ \Psi \simeq \Phi_{i+1}'' * \Phi_i,$$

where

$$\Psi : W(Y_i(\mathbb{S}_i), L_i, \Phi_i^{-1}(\mathbb{K})) \circ W(Y_i, L_i, \mathbb{S}_i) \to W(Y_i, L_i, \varnothing)$$

is a diffeomorphism obtained by handle cancellation, as in Definition 9.6.

(7) *(An index 2/3 birth-death singularity away from Σ): This move is defined analogously to Move* (6)*, but instead we start with a pair of adjacent terms $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ where $\mathbb{S}_i$ is a framed 1-dimensional link, and $\mathbb{S}_{i+1}$ is a framed 2-sphere, which cancels a component of $\mathbb{S}_i$. The pair is replaced with two elementary parametrized cobordisms, $\mathcal{W}'_{i+1}$ and $\mathcal{W}'_i$, with $\mathcal{W}'_{i+1}$ being a trivial cobordism and $\mathcal{W}'_i$ corresponding to surgery on the remaining components of $\mathbb{S}_i$.*

(8) *(A critical value switch between framed 0- or 2-spheres, away from Σ): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ are both of type* (2)*, for 0- or 2-spheres $\mathbb{S}_i \subseteq Y_i$ and $\mathbb{S}_{i+1} \subseteq Y_{i+1}$ such that under the images of the framings of $\mathbb{S}_i$ and $\Phi_i^{-1}(\mathbb{S}_{i+1})$ in Y_i are disjoint, then we replace the pair $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ with a pair*

$$\mathcal{W}'_i = (W'_i, F'_i, \Phi'_i, \mathbb{S}'_i) \qquad \text{and} \qquad \mathcal{W}'_{i+1} = (W'_{i+1}, F'_{i+1}, \Phi'_{i+1}, \mathbb{S}'_{i+1})$$

both of type (2)*, such that $W_i \cup W_{i+1} = W'_i \cup W'_{i+1}$ and $F_i \cup F_{i+1} = F'_i \cup F'_{i+1}$, and also*

$$\Phi_i(\mathbb{S}'_i) = \mathbb{S}_{i+1} \qquad \text{and} \qquad \Phi'_i(\mathbb{S}_i) = \mathbb{S}'_{i+1}.$$

Furthermore, under the canonical identifications

$$W(Y_i(\mathbb{S}_i), L_i, \mathbb{S}'_i) \circ W(Y_i, L_i, \mathbb{S}_i) \cong W(Y_i(\mathbb{S}'_i), L_i, \mathbb{S}_i) \circ W(Y_i, L_i, \mathbb{S}'_i),$$

we have that $\Phi_{i+1} * \Phi_i$ and $\Phi'_{i+1} * \Phi'_i$ are isotopic through diffeomorphisms which fix $Y_i \times \{0\}$.

(9) *(A handleslide amongst the framed 1-dimensional spheres): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ is of type* (2)*, with $\mathbb{S}_i$ a framed 1-dimensional link, then handlesliding a component of $\mathbb{S}_i$ across another component yields a framed link $\mathbb{S}'_i$, and a diffeomorphism $D : W(Y_i, L_i, \mathbb{S}'_i) \to W(Y_i, L_i, \mathbb{S}_i)$ fixing $Y_1 \times \{0\}$, well defined up to isotopy*[6]*. The parametrized cobordism $\mathcal{W}_i$ is replaced by*

$$\mathcal{W}'_i = (W_i, F_i, \Phi_i \circ D, \mathbb{S}'_i).$$

(10) *(A birth-death singularity along $\mathcal{A}$): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \varnothing)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \varnothing)$ are elementary parametrized link cobordisms, both of type* (3)*, such that the two critical arcs $\mathcal{A}_i$ and $\mathcal{A}_{i+1}$ intersect at exactly one point on the common boundary $W_i \cap W_{i+1}$, then we replace the pair $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ with*

$$\mathcal{W}'_i = (W_i \cup W_{i+1}, F_i \cup F_{i+1}, \Phi_{i+1} * \Phi_i, \varnothing),$$

an elementary, parametrized link cobordism of type (1)*.*

(11) *(A critical point switch along $\mathcal{A}$): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \varnothing)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \varnothing)$ are both of type* (3)*, then they are replaced by a pair $\mathcal{W}'_i = (W'_i, F'_i, \Phi'_i, \varnothing)$ and $\mathcal{W}'_{i+1} = (W'_{i+1}, F'_{i+1}, \Phi'_{i+1}, \varnothing)$, such that if a_i and a_{i+1} are the arcs of $\mathcal{A}|_{F_i \cup F_{i+1}}$ which contain the critical arcs of F_i and F_{i+1}, respectively, then $a_i \neq a_{i+1}$ and a_i contains the critical arc of F'_{i+1} and a_{i+1} contains the critical arc of F'_i. We also have that $\Phi_{i+1} * \Phi_i$ and $\Phi'_{i+1} * \Phi'_i$ are isotopic through diffeomorphisms which fix $Y_i \times \{0\}$.*

(12) *(A critical point switch along Σ): If $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ and $\mathcal{W}_{i+1} = (W_{i+1}, F_{i+1}, \Phi_{i+1}, \mathbb{S}_{i+1})$ are two elementary parametrized cobordisms of type* (4)*, then the pair is replaced by a pair*

$$\mathcal{W}'_i = (W'_i, F'_i, \Phi'_i, \mathbb{S}'_i) \qquad \text{and} \qquad \mathcal{W}'_{i+1} = (W'_{i+1}, F'_{i+1}, \Phi'_{i+1}, \mathbb{S}'_{i+1}),$$

both of type (4)*, with $W_i \cup W_{i+1} = W'_i \cup W'_{i+1}$ and $F_i \cup F_{i+1} = F'_i \cup F'_{i+1}$, and*

$$\Phi_i|_{(Y_i \setminus \mathbb{S}_i) \times \{2\}}(\mathbb{S}'_i) = \mathbb{S}_{i+1} \qquad \text{and} \qquad \Phi'_i|_{(Y_i \setminus \mathbb{S}'_i) \times \{2\}}(\mathbb{S}_i) = \mathbb{S}'_{i+1}.$$

Furthermore, under the canonical identifications

$$W(Y_i(\mathbb{S}_i), L_i(\mathbb{S}_i), \mathbb{S}'_i) \circ W(Y_i, L_i, \mathbb{S}_i) \cong W(Y_i(\mathbb{S}'_i), L_i(\mathbb{S}'_i), \mathbb{S}_i) \circ W(Y_i, L_i, \mathbb{S}'_i),$$

[6]An isotopy of an attaching map yields a diffeomorphism of the spaces obtained by attaching a handle, well defined up to isotopy. Handlesliding a handle across another corresponds to an isotopy of the attaching sphere of a handle on the boundary of the 4-manifold obtained by attaching another handle. As such, the diffeomorphism D resulting from handlesliding is well defined up to isotopy.

we have that $\Phi_{i+1}*\Phi_i$ and $\Phi'_{i+1}*\Phi'_i$ are isotopic through diffeomorphisms which fix $Y_i\times\{0\}$, as well as the center points of the two bands.

(13) *(A critical point switch between critical points of $\mathcal{A}$ and Σ): If $\mathcal{W}_i=(W_i,F_i,\Phi_i,\varnothing)$ is of type (3) and $\mathcal{W}_{i+1}=(W_{i+1},F_{i+1},\Phi_{i+1},\mathbb{S}_{i+1})$ is of type (4), then we replace the pair $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ with a pair*

$$\mathcal{W}'_i=(W'_i,F'_i,\Phi'_i,\mathbb{S}'_i),\qquad \text{and}\qquad \mathcal{W}'_{i+1}=(W'_{i+1},F'_{i+1},\Phi'_{i+1},\varnothing),$$

such that $\mathcal{W}'_i$ is of type (4) and $\mathcal{W}'_{i+1}$ is of (3). Furthermore, $W_{i+1}\cup W_i=W'_{i+1}\cup W'_i$ and $F_{i+1}\cup F_i=F'_{i+1}\cup F'_i$, and $\Phi_i(\mathbb{S}'_i)=\mathbb{S}_{i+1}$. Furthermore, under the canonical identification (well defined up to isotopy)

$$W(Y_i,L_i,\mathbb{S}'_i)\circ W(Y_i,L_i,\varnothing)\cong W(Y_i(\mathbb{S}'_i),L_i,\varnothing)\circ W(Y_i,L_i,\mathbb{S}'_i),$$

we have that $\Phi_{i+1}*\Phi_i$ and $\Phi'_{i+1}*\Phi'_i$ are isotopic through diffeomorphisms which fix $(Y_i\times\{0\})\cup\{(0,0)\}$.

(14) *(A critical point of Σ crossing $\mathcal{A}$): If $\mathcal{W}_i=(W_i,F_i,\Phi_i,\varnothing)$ is of type (3.I^+) and $\mathcal{W}_{i+1}=(W_{i+1},F_{i+1},\Phi_{i+1},\mathbb{S}_{i+1})$ is of type (4.$\mathbf{w}$), then $\mathcal{W}_i$ and $\mathcal{W}_{i+1}$ are replaced by a pair*

$$\mathcal{W}'_i=(W'_i,F'_i,\Phi'_i,\varnothing)\qquad \text{and}\qquad \mathcal{W}'_{i+1}=(W'_{i+1},F'_{i+1},\Phi'_{i+1},\mathbb{S}'_{i+1}),$$

where $\mathcal{W}'_i$ is of type (3.II^+) and $\mathcal{W}'_{i+1}$ is of type (4.$\mathbf{z}$). Furthermore $W_{i+1}\cup W_i=W'_{i+1}\cup W'_i$ and $F_{i+1}\cup F_i=F'_{i+1}\cup F'_i$. Finally, we have that $\Phi_i^{-1}(\mathbb{S}_{i+1})$ and $\Phi_i^{-1}(\mathbb{S}'_{i+1})$ are equal. The maps $\Phi_{i+1}*\Phi_i$ and $\Phi'_{i+1}*\Phi'_i$ are isotopic diffeomorphisms, relative $Y_i\times\{0\}$. The isotopy between $\Phi_{i+1}*\Phi_i$ and $\Phi'_{i+1}*\Phi'_i$ does not fix the point $(0,0)$ of the band. Instead the image of $(0,0)$ on the surface $\Sigma_i\cup\Sigma_{i+1}$ crosses a single arc of $\mathcal{A}$, transversely, at a single point (cf. Figure 10.1).

To prove invariance of the link Floer cobordism maps, we will prove the following:

Proposition 10.3. *If (W,F) is a link cobordism with $F=(\Sigma,\mathcal{A})$, and each component of W and each component of Σ have a nonempty incoming and outgoing end, then any two parametrized Kirby decompositions of (W,F) are weakly equivalent.*

The remainder of this section will be devoted to the proof of Proposition 10.3.

10.2. **Spaces of adapted coordinates.** We now prove several useful lemmas about spaces of coordinate functions, which will be necessary for various Morse theoretic arguments to come.

If p is a critical point of a Morse function f of index k, we let $\mathcal{CO}_p(f)$ denote the space of local coordinates in which f takes the form

$$f=f(p)-(x_1^2+\cdots+x_k^2)+(x_{k+1}^2+\cdots+x_n^2)$$

(we give $\mathcal{CO}_p(f)$ the C^∞ topology). Note that by picking a base set of coordinates, the set $\mathcal{CO}_p(f)$ is homeomorphic to the space of germs of self diffeomorphisms of $\mathbb{R}^n$ which map $\mathbf{0}$ to $\mathbf{0}$ and preserve the function $f=f(p)-(x_1^2+\cdots+x_k^2)+(x_{k+1}^2+\cdots+x_n^2)$. We let $\mathcal{D}(f)$ denote this space of germs of diffeomorphisms. Let $L\mathcal{D}(f)\subseteq\mathcal{D}(f)$ denote the subset of linear maps in $\mathcal{D}(f)$. The following fact is well known (cf. [Cer70, pg 168]). We will use this trick repeatedly.

Lemma 10.4. *The space $L\mathcal{D}(f)$ is a deformation retraction of $\mathcal{D}(f)$. Furthermore, $L\mathcal{D}(f)$ is isomorphic to $SO(k,n-k)$.*

Proof. The homotopy is given by linearizing. We define

$$H_t(\phi)(\mathbf{x})=\begin{cases}\frac{1}{t}\phi(t\mathbf{x}), & \text{if } t\neq 0\\ d_{\mathbf{0}}(\phi)(\mathbf{x}), & \text{if } t=0\end{cases}.$$

First note that that if $\phi\in\mathcal{D}(f)$, then $H_t(\phi)\in\mathcal{D}(f)$ for all t by the following easy computation:

$$f(H_t(\phi)(\mathbf{x}))=f(\tfrac{1}{t}\phi(t\mathbf{x}))=\tfrac{1}{t^2}(f(\phi(t\mathbf{x}))=\tfrac{1}{t^2}f(t\mathbf{x})=f(\mathbf{x}).$$

To see that the map H_t is a continuous map $\mathcal{D}(f)\times[0,1]\to\mathcal{D}(f)$, we note that if I is a multi-index, then

$$\left(\frac{\partial^I}{\partial\mathbf{x}^I}H_t(\phi)\right)(\mathbf{x})=t^{|I|-1}\left(\frac{\partial^I}{\partial\mathbf{x}^I}\phi\right)(t\mathbf{x}),$$

from which continuity in the C^∞ topology is straightforward. The space $L\mathcal{D}(f)$ is exactly the set of linear maps which preserve the quadratic form f, which is $SO(k, n-k)$ by definition. □

We prove a lemma about the space of coordinates adapted to certain Morse functions on 4-manifolds with embedded surfaces:

Lemma 10.5. *Let (x, y, w, z) be the standard coordinates on $\mathbb{R}^4$, let Σ denote the subset $\{(x, y, 0, 0)\} \subseteq \mathbb{R}^4$, and let $f = -x^2 + y^2 + w^2 + z^2$. Let $\mathcal{D}_\Sigma(f)$ denote the space of germs of self diffeomorphisms of $\mathbb{R}^4$ mapping $\mathbf{0}$ to $\mathbf{0}$, mapping Σ to itself, and preserving the Morse function f. Then $\mathcal{D}_\Sigma(f)$ deformation retracts onto the subset of linear diffeomorphisms $L\mathcal{D}_\Sigma(f) \subseteq \mathcal{D}_\Sigma(f)$, which is itself canonically isomorphic to $S(O(1,1)\times O(2))$, which has four components. The subset of these which preserve a given orientation of Σ has two components. The set of vector fields $\mathcal{V}$ arising as $\phi_*(v_0)$ for*

$$v_0 = 2\left(-x\frac{\partial}{\partial x} + y\frac{\partial}{\partial y} + w\frac{\partial}{\partial w} + z\frac{\partial}{\partial z}\right)$$

and $\phi \in \mathcal{D}_\Sigma(f)$ is connected.

Proof. The deformation retraction of Lemma 10.4 is easily seen to be through maps which preserve Σ and f, so $L\mathcal{D}_\Sigma(f)$ is a deformation retraction of $\mathcal{D}_\Sigma(f)$. Write an element $M \in L\mathcal{D}_\Sigma(f)$ as

$$M = \begin{pmatrix} A & B \\ 0 & D \end{pmatrix},$$

for 2×2 matrices A, B, and D (here the first component corresponds to the space spanned by $\partial/\partial x$, and $\partial/\partial y$ and the second corresponds to the space spanned by $\partial/\partial w$ and $\partial/\partial z$). Asking that M preserves f is easily seen to imply

$$A \in O(1,1), \qquad B = 0, \qquad \text{and} \qquad D \in O(2).$$

Since M must have positive determinant, we see that A and D must both be positively oriented, or both negatively oriented. Hence $L\mathcal{D}_\Sigma(f)$ has four components. Two of these components will preserve the orientation of Σ, and two will reverse it.

To see that $\mathcal{V}$ is connected, it is sufficient to note that the vector field v_0 itself is preserved by the diffeomorphisms obtained by replacing any number of the coordinate functions τ with $-\tau$. We can go between any components of $\mathcal{D}_\Sigma(f)$ in this manner. □

Lemma 10.6. *Let (x, y, w, z) be the standard coordinates on $\mathbb{R}^4$ and let $\Sigma = \{(x, y, 0, 0)\}$. Let $\mathcal{D}_\Sigma(f_\Sigma)$ denote the space of germs of diffeomorphisms which fix $\mathbf{0}$, preserve Σ, and which preserve $f_\Sigma(x, y) = -x^2 + y^2$ on Σ. Letting $L\mathcal{D}_\Sigma(f_\Sigma)$ denote the subset consisting of linear diffeomorphisms, then we have that $\mathcal{D}_\Sigma(f_\Sigma)$ deformation retracts on $L\mathcal{D}_\Sigma(f_\Sigma)$, and that $L\mathcal{D}_\Sigma(f_\Sigma)$ has four components, though the set $\mathcal{F}$ of Morse functions of the form $f \circ \phi$, where $\phi \in \mathcal{D}_\Sigma(f_\Sigma)$, is connected.*

Proof. We note that if ϕ fixes $\mathbf{0}$, preserves Σ, and preserves f_Σ on Σ, then the linearizations $H_t(\phi)$ will do the same for all t. Hence as before we get a deformation retraction of $\mathcal{D}_\Sigma(f_\Sigma)$ onto $L\mathcal{D}_\Sigma(f_\Sigma)$. Write an element $M \in L\mathcal{D}_\Sigma(f_\Sigma)$ as

$$M = \begin{pmatrix} A & B \\ 0 & D \end{pmatrix},$$

we see that see that B can be anything, and $A \in O(1,1)$. Deformation retracting onto the subset with $B = 0$ we see that $\mathcal{D}_\Sigma(f_\Sigma)$ is homotopy equivalent to $(O(1,1) \times GL(\mathbb{R}^2))^+$, which has four components.

The images of the function $f = -x^2 + y^2 + w^2 + z^2$ under all of these components are easily seen to coincide, since replacing any of the coordinate functions τ with $-\tau$ preserves f. □

We also need to consider spaces of gradient like vector fields which take a standard form near critical points on Σ. We make the following definition:

Definition 10.7. *Suppose that (W, F) is a link cobordism and $F = (\Sigma, \mathcal{A})$ and that f_Σ is a fixed Morse function on Σ, with only index 1 critical points. We define $\mathcal{M}(f_\Sigma)$ to be the space of Morse functions on W which restrict to f_Σ on Σ, and such that*

(1) $\mathrm{Crit}(f_\Sigma) = \mathrm{Crit}(f) \cap \Sigma$;

(2) *There are coordinates (x, y, w, z) (oriented, for W) near each point $p \in \mathrm{Crit}(f_\Sigma)$ such that $f = f(p) - x^2 + y^2 + w^2 + z^2$, and Σ is locally given by $\{(x, y, 0, 0)\}$,*
(3) *All critical points of $f|_{W\setminus\Sigma}$ have distinct critical values.*

*If f_Σ is a fixed Morse function on Σ and $f \in \mathcal{M}(f_\Sigma)$, we call a vector field v a **gradient like vector field** for f if the following hold:*

(1) *$v(f) > 0$ away from $\mathrm{Crit}(f)$;*
(2) *v is tangent to Σ along Σ;*
(3) *near any critical point $p \in \mathrm{Crit}(f) \setminus \Sigma$ of index k, there are coordinates (x_1, x_2, x_3, x_4) such that*

$$f = -x_1^2 - \cdots - x_k^2 + x_{k+1}^2 + \cdots + x_4^2, \quad \text{and} \quad v = 2\left(-x_1\tfrac{\partial}{\partial x_1} - \cdots - x_k\tfrac{\partial}{\partial x_k} + x_{k+1}\tfrac{\partial}{\partial x_k} + \cdots + x_4\tfrac{\partial}{\partial x_4}\right);$$

(4) *Near each $p \in \mathrm{Crit}(f_\Sigma)$, there are coordinates (x, y, w, z), (oriented, for W and Σ) such that Σ is locally given by $\{(x, y, 0, 0)\}$,*

$$f = f(p) - x^2 + y^2 + w^2 + z^2, \qquad \text{and} \qquad v = 2\left(-x\tfrac{\partial}{\partial x} + y\tfrac{\partial}{\partial y} + w\tfrac{\partial}{\partial w} + z\tfrac{\partial}{\partial z}\right).$$

We write $\mathcal{V}(f)$ for the set of gradient like vector fields, in the above sense, for f.

Remark 10.8. If $f \in \mathcal{M}(f_\Sigma)$, the existence of coordinates in a neighborhood of each $p \in \mathrm{Crit}(f_\Sigma)$ as in the above definition imply that p is an index 1 critical point for both f and f_Σ.

Lemma 10.9. *If f_Σ is a fixed Morse function on Σ and $f \in \mathcal{M}(f_\Sigma)$, then the set $\mathcal{V}(f)$ of gradient like vector fields for f is nonempty, and connected.*

Proof. We first show $\mathcal{V}(f)$ is nonempty. First note that at the critical points of f_Σ, we can locally find a gradient like vector field, simply as the gradient under a choice of local coordinates that puts the critical point into the local form in the definition of $\mathcal{M}(f_\Sigma)$. Similarly near a critical point of $f|_{W\setminus\Sigma}$, we can find coordinates which give a gradient like vector field of the correct form. We can now pick a finite set $\mathcal{U}$ of open sets U covering W such that the following hold:

(1) U is one of the coordinate charts for a critical point (either on Σ or away from Σ), constructed above. Such open sets can be chosen to be pairwise disjoint;
(2) U does not intersect Σ or any of the critical points of f;
(3) U has coordinates (x, y, w, z) centered at a point of Σ, such that locally $(x, y, 0, 0)$ corresponds to Σ, and on Σ we have $f|_\Sigma = x$.

On the coordinate charts of the first type, we let associate a vector field v_U, by taking the gradient of f using the local coordinates. On sets of the second type, we let v_U be ∇f, for any Riemannian metric f. On sets of the third type, we let $v_U = \partial/\partial x$. Noting that $v_U(f|_\Sigma) > 0$ for vector fields of this third type, we can shrink the sets of the third type and add sets of the second so that $v_U(f) > 0$ on the entirety of a set U of the third type. We now pick a partition of unity for $\mathcal{U}$, and let v be the associated sum.

We now show that $\mathcal{V}(f)$ is connected. If v and v' are equal on neighborhoods of the critical points of f, then any convex combination will also be a gradient like vector field with a valid local model near each critical point. By Lemma 10.5, we can find a path ϕ_t of local diffeomorphisms of W defined near p, which fix p, preserve f and Σ, and such that $\phi_0 = \mathrm{id}$ and $(\phi_1)_*(v) = v'$. Let ρ be a smooth function which is constantly 1 near p, and has support inside an open set U containing p which is in the domain and image of all of the ϕ_t. We then define

$$v_t|_x = \rho(x) \cdot (\phi_t)_*(v|_{\phi_t^{-1}(x)}) + (1 - \rho(x)) \cdot v|_x. \tag{11}$$

Outside of U, we have $v_t = v$, and hence each v_t is gradient like away from U. Inside of U, but away from p, v_t is pointwise a convex combination of two vector fields V_i^t which each have $V_i^t(f) > 0$, and hence the pointwise convex combination also satisfies that inequality. In a neighborhood of p, $v_t = (\phi_t)_*(v)$, which by assumption takes has the desired local form for the coordinates induced by ϕ_t. Hence v_t is a path of vector fields in $\mathcal{V}(f)$. In this way, we reduced to the case that the two gradient like vector fields are equal in a neighborhood of the critical points, completing the proof by our previous argument. □

10.3. **Changing the Morse function away from Σ.** In this subsection, we construct appropriately nice paths between Morse functions which agree on Σ.

Lemma 10.10. *If f and f' are an arbitrary pair of Morse functions in $\mathcal{M}(f_\Sigma)$, then they can be connected by a path of functions f_t, such that $f_t \in \mathcal{M}(f_\Sigma)$ for all but finitely many t, and at the finite points of time when f_t fails to be in $\mathcal{M}(f_\Sigma)$, one of the following occurs:*

1. *A critical point birth-death occurs away from Σ;*
2. *Two critical points exchange relative value.*

Furthermore, for all t, we have $\mathrm{Crit}(f_\Sigma) = \mathrm{Crit}(f_t) \cap \Sigma$. In addition, for each t, there is a neighborhood of each $p \in \mathrm{Crit}(f_\Sigma)$, with coordinates (x, y, w, z) with $f_t = f_\Sigma(p) - x^2 + y^2 + w^2 + z^2$ where Σ is locally given by $\{(x, y, 0, 0)\}$.

Proof. Let us first prove the claim in the case that f and f' agree in a regular neighborhood of each critical point $p \in \mathrm{Crit}(f_\Sigma)$. In this case, we can simply consider $f_t = tf + (1-t)f'$. Notice that f_t has no critical points along Σ except for those in $\mathrm{Crit}(f_\Sigma)$ since it is constant in t on Σ, and in a neighborhood of $\mathrm{Crit}(f_\Sigma)$. A generic perturbation of f_t near the set of critical points on $W \setminus \Sigma$ will have only birth death singularities, and critical point value switches. Since f_t is already generic near each $p \in \mathrm{Crit}(f_\Sigma)$ (as it's constant near each p), one does not need to perturb near those critical points. On $\Sigma \setminus \mathrm{Crit}(f_\Sigma)$, the functions f_t are constant in t, and also have no critical points since they are equal to f_Σ. Hence no perturbation is needed near Σ.

Using the previous argument, we now just need to show that if f and f' are arbitrary Morse functions in $\mathcal{M}(f_\Sigma)$, then there is a path f_t in $\mathcal{M}(f_\Sigma)$ such that $f_0 = f$ and f_1 agrees with f' in a neighborhood of each point of $\mathrm{Crit}(f_\Sigma)$. Using Lemma 10.6 the space of oriented local diffeomorphisms preserving Σ and f_Σ has four components, but the spaces of Morse functions $-x^2 + y^2 + w^2 + z^2$ induced by these four components are the same. In particular, we can find a path ϕ_t of local diffeomorphisms near a given critical point p, which preserve p, Σ and f_Σ, such that $f' = f \circ \phi_1$, near p. Pick a closed ball B centered at p such that B is in the domain and image of all of the ϕ_t, and let ρ be a bump function which is 1 in a neighborhood of p, and zero outside of B. Consider the time dependent vector field V_t defined by

$$V_t(x) = \rho(x) \cdot \left.\frac{d}{ds}\right|_{s=0} \phi_{t+s}(\phi_t^{-1}(x)).$$

We integrate $V_t(x)$, to define a 1-parameter family Φ_t of diffeomorphisms. We note that Φ_t preserves Σ, f_Σ and p. We note that in a neighborhood of p, we have that Φ_t and ϕ_t agree. Defining $f_t = f \circ \Phi_t$, we have thus connected f by a path in $\mathcal{M}(f_\Sigma)$ to a Morse function which agrees with f' in a neighborhood of p. Doing this for each $p \in \mathrm{Crit}(f_\Sigma)$ reduces the lemma to the previous case, finishing the proof. □

We now define a first class of Morse functions, which we call "nice Morse functions". Since there is no reference to the set of divides, they cannot quite be used to define the link cobordism maps. We will later define a class of Morse functions called "very nice Morse functions", which will also respect the divides, and which can be used to compute the cobordism maps.

Definition 10.11. *If (W, F) is a decorated link cobordism with $F = (\Sigma, \mathcal{A})$, and f_Σ is a Morse function on Σ with only index 1 critical points we say that $f \in \mathcal{M}(f_\Sigma)$ is **nice** if*

1. *All critical points of $f|_\Sigma$, occur before any critical points of $f|_{W\setminus\Sigma}$;*
2. *$f|_{W\setminus\Sigma}$ has only index 1,2 or 3 critical points;*
3. *All index 1 critical points of $f|_{W\setminus\Sigma}$ occur before the index 2 and 3 critical points, and the index 2 critical points all occur before the index 3 critical points.*
4. *All critical points of $f|_{W\setminus\Sigma}$ have distinct critical values.*

We let $\mathcal{M}^{\mathrm{nice}}(f_\Sigma)$ denote the set of nice Morse functions restricting to f_Σ on Σ.

Definition 10.12. *We say that a smooth path of functions f_t is a **nice path** if f_t is in $\mathcal{M}^{\mathrm{nice}}(f_\Sigma)$ for all but finitely many t, and at those points of time, one of the following singularities occurs:*

1. *two critical points of $f|_{W\setminus\Sigma}$ exchange relative value;*
2. *a birth-death singularity occurs, with higher critical value than any of the critical values of of $f|_\Sigma$.*

Remark 10.13. Notice that if f_t is a nice path, then necessarily all birth deaths occur away from Σ, since f_t is constant on Σ. Similarly, if a birth-death singularity occurs, then it necessarily must be an index 1/2 or 2/3 birth-death, since the Morse functions before and after the singularity are nice, by definition.

Lemma 10.14. *Given any two Morse functions $f_0, f_1 \in \mathcal{M}^{\text{nice}}(f_\Sigma)$, there is a nice path f_t from f_0 to f_1. If $\mathcal{A}$ is a set of divides on Σ, and the functions f_0 and f_1 map W into $[0,2]$, all critical values of $f_i|_{W\setminus\Sigma}$ are greater than 1, and all critical values of $f_i|_\Sigma$ and $f_i|_\mathcal{A}$ are less than 1, then we can further assume that all critical values of $f_t|_{W\setminus\Sigma}$ are also greater than 1 for all t.*

Proof. Using Lemma 10.10, we can find a generic path f_t from f_0 to f_1, which is in $\mathcal{M}(f_\Sigma)$ for all but finitely many t, where a critical point value switch occurs, or a birth death singularity occurs. We pick a path of gradient like vector fields v_t for f_t. We note that by assumption v_t is a gradient like vector field such that the descending manifolds of $\text{Crit}(f_\Sigma)$ do not intersect the ascending manifolds of any critical points of $f_t|_{W\setminus\Sigma}$, since all of the descending flowlines from points in $\text{Crit}(f_\Sigma)$ stay in Σ, by assumption. Hence we can modify the path f_t so that all critical points of any $f_t|_{W\setminus\Sigma}$ occur above all of critical points in $\text{Crit}(f_\Sigma)$, using the "independent trajectories principle" ([Kir78, pg. 40]). Indeed we can also assume that they occur above any critical point of $f|_\mathcal{A}$. This can be done while preserving $f_t|_\Sigma$.

Now one can use the techniques of Cerf graphics, developed in [Kir78] (cf. [Cer70] and [Hat73]) to manipulate the Cerf graphic associated to the path f_t so that all index 1 critical points occur before index 2 critical points, which occur before index 3 critical points, and that there are also no index 0 or 4 critical points. There are finitely many points of time when a critical point birth death may occur, or when a critical point value swap may occur. If the the critical values of each $f_t|_{W\setminus\Sigma}$ are greater than 1, then the resulting path (after performing Cerf theory) can also be constructed to satisfy that. □

10.4. **Changing the Morse function f_Σ on Σ.** We now consider different Morse functions f_Σ on Σ. We first define the a set of Morse functions on Σ which respect the divides $\mathcal{A}$ well enough that we can use them to compute the link cobordism maps.

Definition 10.15. *Given a surface with divides $(\Sigma, \mathcal{A})$ with boundary partitioned into subsets $-L_1$ and L_2, we say a Morse function f on Σ is $\mathcal{A}$-**admissible** if the following hold:*

1. *The level sets of f are nonempty;*
2. *Each component of $f^{-1}(t)$ intersects $\mathcal{A}$ nontrivially for each t;*
3. *There are no index 0 or index 2 critical points;*
4. *$f|_\mathcal{A}$ is Morse;*
5. *All critical points of f and $f|_\mathcal{A}$ have distinct values;*

Remark 10.16. If f is an $\mathcal{A}$-admissible Morse function on Σ, then the critical points of f cannot intersect $\mathcal{A}$, since such a critical point would be a critical point of f and of $f|_\mathcal{A}$, contradicting condition (5) of the previous definition.

Lemma 10.17. *If $(\Sigma, \mathcal{A})$ is a surface with divides with boundary partitioned into incoming and outgoing components, such that each component of Σ has nonempty incoming and outgoing boundary components, then there exists an $\mathcal{A}$-admissible Morse function on Σ, and any two $\mathcal{A}$-admissible Morse functions, f_0 and f_1, on Σ can be connected by a path of Morse functions f_t such that there are only finitely many points of time t where f_t is not $\mathcal{A}$-admissible. At those points of time, one of the following singularities occurs:*

1. *A birth-death singularity of $f_t|_\mathcal{A}$ occurs.*
2. *Two critical points of $f_t|_\mathcal{A}$ switch relative ordering.*
3. *Two critical points of f_t exchange relative ordering.*
4. *A critical point of $f_t|_\mathcal{A}$ and a critical point of f exchange relative ordering.*
5. *A critical point of f_t crosses $\mathcal{A}$, transversely.*

Each of the above events can be assumed to occur at distinct times t.

Proof. Suppose that f_0 and f_1 are two $\mathcal{A}$-admissible Morse functions. Using standard Morse theory, one can find a path of functions f_t which are Morse on Σ except at finitely many t, where a birth/death singularity occurs, and such that there are no index 0 or index 2 critical points, throughout (cf. [Juh15b, Theorem 2.21]). Since there are no index 0 or index 2 critical points of f_t for any t, there can be no birth/death singularities

at all, so f_t is Morse for all t. To ensure that each component of $(f_t)^{-1}(s)$ intersects $\mathcal{A}$ nontrivially, one can then change the path f_t of Morse functions by first taking f_0 and considering the path of Morse functions $f_0 \circ \phi_t$, where ϕ_t is an isotopy of Σ which "scrambles" the surface suitably, concatenating with the path of Morse functions $f_t \circ \phi_1$, then concatenating with the path of Morse functions $f_1 \circ \phi_{1-t}$. Such a path has the property that $(f_t)^{-1}(s)$ intersects $\mathcal{A}$ nontrivially for each t and s. After performing a generic perturbation of f_t, for all but finitely many t, f_t will be Morse when restricted to $\mathcal{A}$, all critical values of f_t and $f_t|_{\mathcal{A}}$ will be distinct and occur at distinct points on Σ. At finitely many t, one of the singularities listed can occur. □

Definition 10.18. *If $(W, \Sigma, \mathcal{A})$ is a decorated link cobordism, we will say that a Morse function $f : W \to [0,2]$ is **very nice**, and write $f \in \mathcal{M}^{\text{nice}}(\mathcal{A})$, if the following holds:*

(1) f is a nice Morse function;
(2) $f|_\Sigma$ is $\mathcal{A}$-admissible;
(3) All critical values of $f|_\Sigma$ or $f|_{\mathcal{A}}$ are less than 1, and all critical values of $f|_{W\setminus\Sigma}$ are greater than 1;
(4) All critical values of f and $f|_{\mathcal{A}}$ are distinct.

*We say that a path of smooth functions $f_t : W \to [0,2]$ is a **very nice path** if each component of $(f_t)^{-1}(s)$ intersects $\mathcal{A}$ nontrivially for each s and t, f_t is a nice path, and each f_t is a very nice Morse function for all but finitely many t, where f_t fails to be a nice Morse function, or one of the singularities listed in Lemma 10.17 occurs. Furthermore, all such singularities occur at distinct times t.*

We now prove the following:

Lemma 10.19. *Suppose (W, F) is a link cobordism and $(f_\Sigma)_0$ and $(f_\Sigma)_1$ are two $\mathcal{A}$-admissible Morse functions from Σ to $[0,2]$ such that all critical points of $(f_\Sigma)_0, (f_\Sigma)_1, (f_\Sigma)_0|_{\mathcal{A}}$ and $(f_\Sigma)_1|_{\mathcal{A}}$ have value less than 1. Then there is a very nice path f_t of smooth functions from f_0 to f_1 such that $(f_0)|_\Sigma = (f_\Sigma)_0$ and $(f_1)|_\Sigma = (f_\Sigma)_1$.*

Our proof is by explicit construction. We first describe the construction, and then prove that it can be used to construct the very nice path in the previous lemma.

Equip a neighborhood of $Y_1 \cup \Sigma$ with a Riemannian metric. This yields an identification of $N_{r_0}(\Sigma)$, the disk bundle of radius r_0 inside of the normal bundle, with a regular neighborhood of Σ in W. Let

$$\pi : N_{r_0}(\Sigma) \to \Sigma,$$

denote projection of the normal bundle onto Σ. Suppose $f_\Sigma : \Sigma \to [0,2]$ is a Morse function which has only index 1 critical points, and all critical values of f_Σ and $f_\Sigma|_{\mathcal{A}}$ are strictly less than 1. Writing Σ_1 for $f_\Sigma^{-1}([0,1])$ we assume that $N_2(\Sigma_1)$ (the radius 2 closed disk bundle) intersects Y_2 (the outgoing boundary of W) trivially. Pick a Morse function $h : N(Y_1) \to [0,2]$ on a collar neighborhood of Y_1 such that $h = k\cdot(f_\Sigma \circ \pi)$ for a constant $k > 1$, on the intersection of $N_{r_0}(\Sigma)$ and $N(Y_1)$. We assume that $N(Y_1) \cap Y_2 = \varnothing$ as well. We also assume that h and $N(Y_1)$ have been selected so that f_Σ and $f_\Sigma|_{\mathcal{A}}$ have no critical points on $N(Y_1)$. Let

$$r : N_{r_0}(\Sigma) \to [0, r_0]$$

denote the radial distance function. For notational simplicity, normalize the metric so that $r_0 = 2$. We will occasionally write

$$N_s(Y_1) \subseteq N(Y_1)$$

for the set of points x where $h(x) \leq s$.

Let

$$\rho : [0,2] \to [0,1]$$

be a smooth function such that the following hold:

(1) $0 \leq \rho \leq 1$;
(2) ρ is nondecreasing, and $\rho'(t) > 0$ at any t where $\rho(t) \neq 1$;
(3) $\rho(t) = t$ for t in a neighborhood of 0, and ρ is constantly 1 in a neighborhood of $t = 2$.

Let $F : N(Y_1) \cup N_2(\Sigma) \to \mathbb{R}$ be defined by $f_\Sigma \circ \pi$, or h/k, whichever is defined (by assumption they are equal when they are both defined).

We now consider the Morse function

$$f = F + (\rho \circ r)^2(\rho \circ h), \tag{12}$$

viewed as a function on $N(Y_1) \cup N_2(\Sigma)$. We define

$$W_0^f = f^{-1}([0,1]).$$

Lemma 10.20. *The function f is Morse on $W_0^f = f^{-1}([0,1])$, a neighborhood of $Y_1 \cup \Sigma_1$. The critical points of f coincide exactly with the critical points of f_Σ, and each critical point of f_Σ (which is of index 1 by assumption) corresponds to an index 1 critical point of f. There are local coordinates (x, y, w, z) near any critical point p on Σ such that Σ is locally given by $\{(x, y, 0, 0)\}$ and*

$$f = f_\Sigma(p) - x^2 + y^2 + w^2 + z^2.$$

Furthermore, W_0^f is a cobordism from Y_1 to the closed, properly embedded submanifold, $f^{-1}(1)$, of W.

Proof. Note first that if $\rho(r(x)) = \rho(h(x)) = 1$, then $f(x) > 1$, so $x \notin W_0^f$. Hence if $x \in W_0^f$, at least one of $\rho(r)$ and $\rho(h)$ must be strictly less than 1. If $0 < \rho(r(x)) < 1$, then $x \in N_2(\Sigma)$, and we can compute

$$\frac{d}{dr} f = 2\rho(r)\rho'(r)\rho(h),$$

which is strictly positive when $0 < \rho(r) < 1$ and $h > 0$. Hence there are no critical points with $0 < \rho(r) < 1$. If $\rho(r) = 1$ at a point $x \in W_0^f$, then $\rho(h) < 1$ at x, by what we have said, so $x \in N_2(Y_1)$. Hence $f = h/k + \rho(h)$, which is easily seen to not have any critical points on $N_2(Y_1)$.

Hence the only remaining case where critical points can occur is $\rho(r) = 0$, which corresponds exactly to being on Σ. Since $f|_\Sigma = f_\Sigma$, critical points of f can only occur at critical points of f_Σ. Noting f_Σ has no critical points when $\rho(h) < 1$, we see that at any critical point of f_Σ, we have $\rho(h) = 1$. In particular, near any critical point for f, we have that locally,

$$f = (f_\Sigma \circ \pi) + r^2.$$

Picking coordinates (x, y) on Σ near a critical point p of f_Σ, such that $f_\Sigma = -x^2 + y^2$, and extending these to coordinates (x, y, w, z) near p on the $SO(2)$ disk bundle $N_2(\Sigma)$ such that $r^2 = w^2 + z^2$, we have

$$f = f_\Sigma(p) - x^2 + y^2 + w^2 + z^2.$$

We now need to show that $Y' = f^{-1}(1)$ is a closed, properly embedded submanifold of W, and that W_0^f is a cobordism from Y_1 to Y'. The claim that it is properly embedded is equivalent to showing that it doesn't intersect ∂W. Note that $Y_1 = f^{-1}(0)$, so $Y_1 \cap Y' = \varnothing$. We note $W_0^f \subseteq N_2(Y_1) \cup N_2(\Sigma_1)$ (where $\Sigma_1 = f_\Sigma^{-1}([0,1])$), and the latter two sets intersect Y_2 trivially by assumption. Hence Y' is a compact, embedded submanifold in $\operatorname{int} W$ and $W_0^f \subseteq W$ is a cobordism from Y_1 to Y'. □

As an intermediate technical results, we now prove a very slight modification of [GWW, Lemma 3.1].

Lemma 10.21. *Suppose that $f_t : W \to [0,1]$ is a path of Morse functions on (W, Σ) such that $f_t \in \mathcal{M}(f_t|_\Sigma)$ for all t (in particular all critical values are distinct, and there are no birth death singularities of f_t, but we are assuming nothing about a dividing set). Then there is a smooth isotopy $\psi_t : W \to W$ such that $\psi_t|_{\partial W} = \mathrm{id}$ for all t, and an isotopy $\phi_t : \mathbb{R} \to \mathbb{R}$ such that*

$$f_t = \phi_t \circ f_0 \circ \psi_t^{-1}.$$

Proof. If we don't require ψ_t to map Σ to Σ, then [GWW, Lemma 3.1] gives us the result immediately. We briefly sketch the modification, to our setting. First pick a $\phi_t : \mathbb{R} \to \mathbb{R}$ such that $\phi_t \circ f_t$ and f_0 have the same critical values. Such a ϕ_t exists because by assumption all critical values of f_t are distinct, and there are no critical point birth death singularities. Write g_t for $\phi_t \circ f_t$. Pick a path of gradient like vector fields V_t for g_t such that $V_t|_\Sigma$ is tangent to Σ. Away from Σ and the critical points of g_t, we define a vector field

$$\overline{V}_t = -\frac{\partial_t g_t}{V_t(g_t)} \cdot V_t.$$

Note that

$$dg_t(\overline{V}_t) + \partial_t g_t = 0.$$

Extend $\overline{V}_t$ over any critical points p of g_t by picking local coordinates $\theta_t : B^4 \to W$ such that, if $p \in \Sigma$, we have

$$(g_t \circ \theta_t)(x, y, w, z) = c - x^2 + y^2 + w^2 + z^2$$

(with Σ locally given by $\{(x, y, 0, 0)\}$), and in a neighborhood of the critical point p defining

$$\overline{V}_t(y) = \frac{d}{ds}\Big|_{s=0} \theta_{s+t}(\theta_t^{-1}(y)).$$

We use a similar formula near critical points which are not on Σ. We patch together the various definition of $\overline{V}_t$ using partitions of unity, and note that they satisfy

$$dg_t(\overline{V}_t) + \partial_t g_t = 0.$$

We define ψ_t to be the flow of the time dependent vector field $\overline{V}_t$, and note that the latter differential equation implies $g_t \circ \psi_t$ is constant, implying the claim. □

We can now prove Lemma 10.19:

Proof of Lemma 10.19. Suppose $(f_\Sigma)_0 : \Sigma \to [0, 2]$ and $(f_\Sigma)_1 : \Sigma \to [0, 2]$ are two $\mathcal{A}$-admissible Morse functions, such that all critical values of $(f_\Sigma)_0, (f_\Sigma)_1, (f_\Sigma)_0|_{\mathcal{A}}$ and $(f_\Sigma)_1|_{\mathcal{A}}$ are less than 1. Using Lemma 10.17, we can find a path $(f_\Sigma)_t$ of Morse functions on Σ from $(f_\Sigma)_0$ to $(f_\Sigma)_1$, which are $\mathcal{A}$-admissible for all but finitely many t, where one of the listed singularities occurs. We arrange so that all critical values of $(f_\Sigma)_t$, and $(f_\Sigma)_t|_{\mathcal{A}}$ occur inside of the interval $(0, 1)$, for all t. We then use the construction in Equation (12) to extend each $(f_\Sigma)_t$ smoothly to a smooth function $f_{t,0}$ which is defined on neighborhoods of $Y_1 \cup \Sigma$. We now extend $f_{0,0}$ to a Morse function f_0 on all of W with $f_0|_\Sigma = (f_\Sigma)_0$. Recalling the notation $W_0^{f_{0,0}} = f_{0,0}^{-1}([0, 1])$, since each component of $W \setminus W_0^{f_{0,0}}$ has a nonempty incoming and outgoing end, and $\Sigma \cap (W \setminus W_0^{f_{0,0}})$ is diffeomorphic to a collection of cylinders with trivial divides, we can pick the extension f_0 on W such that f_0 has only critical points of index 1,2 and 3, all critical values are distinct, and all index 1 critical points occur before index 2 and 3 critical points, and index 2 critical points occur before index 3 critical points. Furthermore, we can assume that on $W \setminus W_0^{f_{0,0}}$, the functions $f_0|_\Sigma$ and $f_0|_{\mathcal{A}}$ have no critical points.

By adapting the proof of Lemma 10.21, we can find a 1-parameter family of diffeomorphisms

$$\psi_t : W \setminus W_0^{f_{0,0}} \to W \setminus W_0^{f_{t,0}},$$

mapping $\Sigma \setminus W_0^{f_{0,0}}$ to $\Sigma \setminus W_0^{f_{t,0}}$ such that $(f_\Sigma)_t \circ \psi_t = (f_\Sigma)_0$. Using this family ϕ_t we can simply push forward the Morse function f_0 on $W \setminus W_0^{f_{0,0}}$ to smoothly extend the family $f_{t,0}$ to a family of smooth functions f_t on all of W, which restrict to $f_{t,0}$ on $W_0^{f_{t,0}}$. This causes no new birth death singularities and no critical point value switches away from Σ, since $f_t|_{W \setminus W_0^{f_{t,0}}}$ is simply the push forward of $f_{t,0}$ to $W \setminus W_0^{f_{t,0}}$. Also all critical values of $f_t|_\Sigma$ and $f_t|_{\mathcal{A}}$ are less than 1, and all critical values of $f_t|_{W \setminus \Sigma}$ are greater than 1. At the finitely many points where $(f_\Sigma)_t$ is not $\mathcal{A}$-admissible, the functions f_t are nice but fail to be very nice, and at each such t one of the listed moves occurs. □

Combining Lemmas 10.14 and 10.19, we get the following:

Lemma 10.22. *Suppose $(W, \Sigma, \mathcal{A})$ is a decorated link cobordism, and that f_0 and f_1 are two very nice Morse functions from W to $[0, 2]$. Then there is smooth path of functions $f_t : W \to [0, 2]$ which is a very nice path.*

Proof. Using Lemma 10.19, we can find a very nice path of Morse functions f_t' such that $f_0'|_\Sigma = f_0|_\Sigma$ and $f_1'|_\Sigma = f_1|_\Sigma$. Now by Lemma 10.14, we can find nice paths of functions from f_0' to f_0 and a nice path of functions from f_1' to f_1. By definition these are paths of functions which are constant along Σ (and hence their restrictions to Σ are $\mathcal{A}$-admissible), and are nice Morse functions, except at finitely many t, when a critical value swap away from Σ, or a birth-death singularity away from Σ, may occur. Furthermore, by the second part of Lemma 10.14, we can also assume that the critical values of the two paths restricted to $W \setminus \Sigma$ are all in the interval $(1, 2)$, and that the critical values of the functions restricted to Σ and $\mathcal{A}$ all lie in $(0, 1)$.

Concatenating these three paths, we get a very nice path of functions from f_0 to f_1. □

10.5. **Morse data and parametrized Kirby decompositions.** We now describe how a collection of Morse data on a decorated link cobordism gives a parametrized Kirby decomposition, well defined up to weak equivalence.

Definition 10.23. *Suppose that $(W, \Sigma, \mathcal{A})$ is a decorated link cobordism and that $f : W \to [0,2]$ is a Morse function and $\mathbf{b} \subseteq [0,2]$ is a collection of regular values of f, with $\mathbf{b} = (b_1, \dots, b_n)$ with $0 = b_1 < \cdots < b_n = 2$. Define $W_i = f^{-1}([b_i, b_{i+1}])$ and $\Sigma_i = \Sigma \cap W_i$. We say that $(f, \mathbf{b})$ is **very nice Morse data** if f is a very nice Morse function and on each (W_i, Σ_i), exactly one of the following holds*

1. *f and $f|_{\mathcal{A}}$ have no critical points in W_i (i.e. $(W_i, \Sigma_i, \mathcal{A}_i)$ is a trivial cobordism);*
2. *$f|_{W \setminus \Sigma}$ has a single critical point in W_i, which is of index 1 or 3;*
3. *$f|_{W \setminus \Sigma}$ has (arbitrarily many) index 2 critical points in W_i, and no other W_i contains any index 2 critical points;*
4. *$f|_{\Sigma}$ has a single index 1 critical point in W_i (which by assumption is also an index 1 critical point of f);*
5. *$f|_{\mathcal{A}}$ has a single critical point in W_i.*

To a collection of very nice Morse data, $(f, \mathbf{b})$, on a decorated link cobordism $(W, \Sigma, \mathcal{A})$ we now describe a parametrized Kirby decomposition

$$\mathcal{K}(f, \mathbf{b}).$$

In Lemma 10.24 we show that $\mathcal{K}(f, \mathbf{b})$ is determined by f and $\mathbf{b}$ up to weak equivalence.

Write $W_i = f^{-1}([b_i, b_{i+1}])$, $\Sigma_i = \Sigma \cap W_i$ and $\mathcal{A}_i = \Sigma_i \cap \mathcal{A}$. Also define $F_i = (\Sigma_i, \mathcal{A}_i)$. Each (W_i, Σ_i) is an (unparameterized) elementary cobordism in the sense of Definition 9.3, with a Morse function $f_i = f|_{W_i}$, which restricts to be Morse on Σ_i and $\mathcal{A}_i$. To construct a parametrized Kirby decomposition, we need to construct a parametrization on each (W_i, Σ_i). Using Lemma 10.9, we can find a gradient like vector field $v \in \mathcal{V}(f)$, which restricts to Σ to be a gradient like vector field of $f_i|_{\Sigma_i}$.

Suppose first that the Morse function f_i on (W_i, Σ_i) has no critical points. We set $\mathbb{S}_i = \varnothing$, and the flow of $v/v(f)$ immediately gives a parameterizing diffeomorphism Φ_i. Also note that by Lemma 10.9, the space of gradient like vector fields for f_i is connected, so any two parametrizations Φ_i made in this way are isotopic relative Y_i. The parametrized cobordism $(W_i, F_i, \Phi_i, \varnothing)$ is then an elementary parametrized cobordism of type (1) or (3), depending on whether $f_i|_{\mathcal{A}}$ has a critical point or not.

Now suppose that $f_i|_{W_i \setminus \Sigma_i}$ has a single index 1 or index 3 critical point, or arbitrarily many index 2 critical points. We will construct and framed sphere or framed link and a parameterizing diffeomorphism for (W_i, F_i) turning it into an elementary parametrized cobordism of type (2). Suppose first that f_i has a single critical point of index 1, 2 or 3. The vector field v and a choice of local coordinates in a neighborhood of the critical point give a well defined parametrization. One description of this procedure can be found in [Juh15b, Section 2.3.1], though we will describe an approach that is slightly different, and more obviously gives a smooth parametrization. This approach is described in detail in [Pal63, Section 11]. A similar account, using "spherical modification" instead of handle attachment, can be found in [Kos93, Section VII.2].

Let p be an index k critical point of $f|_{W_i}$ for $k \in \{1,2,3\}$, away from Σ, and suppose that p is the only critical point of f inside of W_i. Let $Y_i = f^{-1}(b_i)$ and $Y_{i+1} = f^{-1}(b_{i+1})$, and define L_i and L_{i+1} as the intersections of Y_i and Y_{i+1} with Σ. For notational convenience, we will describe the handle attachment procedure for arbitrary ambient dimension n (though we will always be assuming that $n = 2$ or $n = 4$). Pick coordinates $(x_1, \dots, x_n)$ near p such that $f = f(p) - (x_1^2 + \cdots + x_k)^2 + (x_{k+1}^2 + \cdots + x_n^2)$. Pick a gradient like vector field v for f on W_i as well, which is equal to

$$v = 2\left(-x_1 \frac{\partial}{\partial x_1} - \cdots - x_k \frac{\partial}{\partial x_k} + x_{k+1} \frac{\partial}{\partial x_k} + \cdots + x_n \frac{\partial}{\partial x_n}\right)$$

near p. The descending manifold is $\{(x_1, \dots x_k, 0, \dots, 0)\}$ and the ascending manifold is $\{(0, \dots, 0, x_{k+1}, \dots, x_n)\}$. Fix a small $\epsilon > 0$, such that the ball of radius 2ϵ is contained in the domain of the coordinates centered at p.

We pick a smooth function $\lambda : \mathbb{R} \to \mathbb{R}$ which has $\lambda'(t) \leq 0$ and satisfies $\lambda(x) = 1$ for $x \leq \frac{1}{2}$, $\lambda(x) > 0$ if $x < 1$ and $\lambda(x) = 0$ if $x \geq 1$. Notice that the space of such functions λ is contractible, since convex

combinations of such functions satisfy the those properties as well. We consider the function

$$g(x) = f(x) - \frac{3\epsilon}{2}\lambda\left(\frac{x_{k+1}^2 + \cdots + x_n^2}{\epsilon}\right).$$

One notes that

$$f^{-1}([b_i, c-\epsilon]) \subseteq g^{-1}([b_i, c-\epsilon]).$$

Write

$$H = g^{-1}([b_i, c-\epsilon]) \setminus f^{-1}([b_i, c-\epsilon]).$$

According to [Pal63, pg. 316], H is homeomorphic to $D^k \times D^{n-k}$, via a homeomorphism which can be canonically specified in terms of λ, ϵ and the choice of coordinates. Furthermore this homeomorphism restricts to a diffeomorphism on $(\text{int } D^k) \times D^{n-k}$, and also restricts to a smooth embedding of $S^{k-1} \times D^{n-k}$ into $f^{-1}(c-\epsilon)$, which is a framed k-sphere in $f^{-1}(c-\epsilon)$. The normalized gradient like vector field, $v/v(f)$, induces a well defined diffeomorphism $Y_i \times [0,1]$ with $f^{-1}([b_i, c-\epsilon])$. The framed k-sphere in $f^{-1}([b_i, c-\epsilon])$ thus naturally induces a framed k-sphere $\mathbb{S}_i$ in $Y_i \times \{0\}$, by using this identification. The homeomorphism of H with $D^k \times D^{n-k}$ described above thus yields a diffeomorphism of $(Y_i \times [0,1]) \cup_{\mathbb{S}_i} (D^k \times D^{n-k})$ with $g^{-1}([b_i, c-\epsilon])$, where the smooth structure on the former space is obtained by attaching the handle and then smoothing corners. By the Proposition from [Pal63, page 319], the gradient like vector field v is transverse to the level sets of g, except at the single critical point p of g (which is also the critical point of f). Hence by flowing along v, the set $f^{-1}([b_i, b_{i+1}]) \setminus g^{-1}([b_i, c-\epsilon])$ is diffeomorphic to $g^{-1}(c-\epsilon) \times [0,1]$, by a diffeomorphism which is unique up to isotopy, relative $g^{-1}(c-\epsilon)$. By picking a collar neighborhood (which is determined up to isotopy) of $g^{-1}(c-\epsilon)$. We thus get a diffeomorphism from $(Y_i \times [0,1]) \cup_{\mathbb{S}_i} (D^k \times D^{n-k})$ to W which is determined up to isotopies fixing $Y_i \times \{0\}$.

Note that even for fixed gradient like vector field and coordinates near p, the construction still relies on a choice of ϵ and λ, though the space of such λ is contractible, as we've already remarked, so any two choices yield framed spheres which are isotopic away from L_i, resulting in Move (2).

In the case that (W_i, F_i) is an elementary cobordism with only index 2 critical points, one can still use the above strategy to arrive at a parametrization $\Phi_i : W(Y_i, L_i, \mathbb{S}_i) \to (W_i, F_i)$. The space of gradient like vector fields for $f|_{W_i}$ is still connected. A generic gradient like vector field will have ascending and descending manifolds of the critical points which are transverse, and hence generically the intersection of the ascending and descending manifolds of two index 2 critical points will be empty, allowing us to perform the above construction whenever the choice of gradient like vector field is generic (paths of gradient like vector fields may of course result in handleslides, as we discuss in the following lemma).

We lastly consider cobordisms of type (4), where the critical point p is of index 1, and is along Σ. In this case, we pick a gradient like vector field v such that in oriented coordinates (x, y, w, z) near the critical point

$$f = f(p) - x^2 + y^2 + w^2 + z^2, \qquad \Sigma = \{(x, y, 0, 0)\}, \qquad \text{and} \qquad v = 2\left(-x\frac{\partial}{\partial x} + y\frac{\partial}{\partial y} + w\frac{\partial}{\partial w} + z\frac{\partial}{\partial z}\right).$$

The space of such coordinates where f takes this form has four connected components by Lemma 10.5. The space of coordinates which restrict to oriented coordinates on Σ has two components. Using one of these sets of coordinates, we can now adapt the above procedure from [Pal63], as we did for critical points which were away from Σ. In this case, using function g constructed above, the procedure gives a parametrization of (W, Σ) as $(Y_i \times [0,1], L_i \times [0,1]) \cup H \cup (g^{-1}(c+\epsilon) \times [1,2])$. Furthermore H is canonically identified as $D^1 \times D^3$. Our choice of local coordinates identified Σ with the subset $\{(x, y, 0, 0)\} \subseteq \mathbb{R}^4$. It is straightforward to check that using this identification, the set $B \subseteq D^1 \times D^3$ defined by $B = \{(x, y, 0, 0) : -1 \leq x \leq 1, -1 \leq y \leq 1\}$ is mapped to Σ, and the identification also sends $(0, 0) \in B$ to the critical point of f.

The construction of the parametrized Kirby decomposition $\mathcal{K}(f, \mathbf{b})$ involved choices of local coordinates, and gradient like vector fields. Nevertheless, we have the following:

Lemma 10.24. *Given a Morse function f and a collection of regular values $\mathbf{b}$ as in Definition 10.23, any two choices of gradient like vector fields and local coordinates in a neighborhood of the critical points of f yield weakly equivalent parametrized Kirby decompositions $\mathcal{K}(f, \mathbf{b})$.*

Proof. We first consider elementary cobordisms of type (1) or (2). We note that the spaces of local coordinates in a neighborhood of each critical point which put f into the correct local model are not connected, though

the components are related by changing the sign of two of the coordinate functions appropriately, by Lemma 10.4. The resulting vector fields are unaffected, and hence the associated spaces of vector fields are connected. By performing Move (5) to change the orientation of the framed sphere $\mathbb{S}_i$, we can assume that the two choices of local coordinates are in the same component in the appropriate space of local coordinates.

On an elementary parametrized cobordism with a Morse function with no critical points, or exactly one critical point, for two choices of local coordinates and associated gradient like vector fields with the correct local model, if the coordinates are in the same component in the appropriate space of local coordinates, then we can simultaneously change the coordinates and the gradient like vector field (cf. Equation (11)). Such a move results in an isotopy of the framed sphere away from L_i (Move (2)).

We now consider the case that f_i has multiple index 2 critical points, away from L_i. In this case, the descending and ascending manifolds of a generic gradient like vector field of each critical point will be transverse, and hence disjoint. For a generic path of gradient like vector fields, there will be finitely many points of time when there is a single flow line between two critical points. At such a point of time, a handleslide of one of the framed 1-spheres occurs, which results in Move (9). Over the course of time when the path of gradient like vector fields have no flow lines amongst each other, the associated parametrized Kirby decompositions are related to each other again by a sequence of Moves (2) and (3).

We lastly consider elementary cobordisms of type (4) (framed 0-spheres along L). By Lemmas 10.5 and 10.9, different choices of local coordinates or gradient like vector field yield framed spheres and parametrizations which are related by an isotopy of the framed sphere along L_i (Move (3)) or a switch of the orientation of the parametrization (Move (5)). □

Lemma 10.25. *Given a parametrized Kirby decomposition $\mathcal{K}$ of (W, F), one can find very nice Morse data $(f, \mathbf{b})$ for (W, F) such that*

$$\mathcal{K} = \mathcal{K}(f, \mathbf{b}).$$

Proof. One can modify the construction from [Juh15b, Lemma 2.14] to get a Morse function on each elementary cobordism which induce the correct parametrization for a choice of gradient like vector field. By modifying them near the boundary, one can ensure that they glue together to form a smooth function on all of W. □

10.6. **Proof of Proposition 10.3.** We now have the tools to prove that all parametrized Kirby decompositions of a given (W, F) are weakly equivalent:

Proof of Proposition 10.3. Given two parametrized Kirby decompositions, $\mathcal{K}_1$ and $\mathcal{K}_2$, one uses Lemma 10.25 to find Morse functions f_1 and f_2, together with collections of regular values, $\mathbf{b}_1$ and $\mathbf{b}_2$, and choices of local coordinates and gradient like vector fields so that

$$\mathcal{K}_1 = \mathcal{K}(f_1, \mathbf{b}_1), \qquad \text{and} \qquad \mathcal{K}_2 = \mathcal{K}(f_2, \mathbf{b}_2).$$

Using Lemma 10.22, we can find a very nice path f_t between f_1 and f_2.

At all but finitely many t, the functions f_t are very nice Morse functions, and hence for some choice of regular values $\mathbf{b}_t$ yield very nice Morse data $(f_t, \mathbf{b}_t)$, which determine a parametrized Kirby decomposition $\mathcal{K}(f_t, \mathbf{b}_t)$ for a choice of gradient like vector field v_t and local coordinates near each critical point. By Lemma 10.24, changing the choices of local coordinates or gradient like vector fields v_t results in a weak equivalence. Changing the set of regular values $\mathbf{b}$ can be achieved by adding and removing points, and hence by applications of Move (4).

At finitely many points of time, a critical point value switch or a critical point birth-death occurs for the functions $f_t|_{W\setminus\Sigma}$. The birth death singularities are limited to index 1/2 and index 2/3 birth death singularities, since nice Morse functions cannot have index 0 or index 4 critical points. At such points of time, by perturbing the path of functions slightly, we can assume that for a small subinterval of time around the singularity, the path of functions f_t is supported in a small ball containing the critical points. The existence of such normal forms is established in [Cer70] (cf. [Juh15b, pg. 17]). As such the parametrized Kirby decompositions immediately preceding and following the singularity may be related by Moves (6), (7) and (8).

If $[a, b]$ is an interval of time which contains no critical point birth-deaths or critical value swaps of f_t, then by Lemma 10.21, we can find an isotopy $\phi_t : \mathbb{R} \to \mathbb{R}$ (ranging over $t \in [a, b]$, with $\psi_a = \mathrm{id}$) and an isotopy

$\psi_t : (W, \Sigma) \to (W, \Sigma)$, which is the identity on ∂W for all t, such that $f_t = \phi_t \circ f_a \circ \psi_t$. In particular, given a choice of $\mathbf{b}_a$, we can take the parametrized Kirby decomposition $\mathcal{K}(f_a, \mathbf{b}_a)$, and push it forward along the diffeomorphism ψ_t to get a parametrized decomposition (of undecorated link cobordisms) $(\psi_t)_*\mathcal{K}(f_a, \mathbf{b}_a)$. A description of the pushforward decomposition is described immediately before Definition 10.1.

For all but finitely many $t \in [a, b]$, the critical values of $f_t|_{\mathcal{A}}$ will be disjoint from the critical values of f_t. Additionally, at all but finitely many t, the parametrization $(\psi_t)_*\Phi_i$ for an embedded 0-sphere along a link will map the point $(0, 0) \in B$ into $\Sigma \setminus \mathcal{A}$. If there are no such t where $f_t|_{\mathcal{A}}$ and f_t share a critical value, and there are no t where the point $(0, 0)$ in a band is mapped to $\mathcal{A}$, then $\mathcal{K}(f_a, \mathbf{b}_a)$ and $\mathcal{K}(f_b, \mathbf{b}_b)$ are related by the strong isotopy move (Move (1)) by Lemma 10.21, where $\mathbf{b}_b = \phi^{-1}(\mathbf{b}_a)$. More generally, at finitely many t in $[a, b]$ one of the singularities listed in Lemma 10.19 may occur. By simply adding and removing points to $\mathbf{b}$ at intermediate steps of the isotopy, we can find a $\mathbf{b}_b$ such that $\mathcal{K}(f_a, \mathbf{b}_a)$ and $\mathcal{K}(f_b, \mathbf{b}_b)$ are related by a sequence of strong isotopies (Move (1)), critical point value switch between a critical point of $f_t|_\Sigma$ and a critical point of $f_t|_{\mathcal{A}}$ (Move (13)), critical point value switches between critical points of $f|_{\mathcal{A}}$ (Move (11)), critical point birth-deaths of $f_t|_{\mathcal{A}}$ (Move (10)), or isotopies of a parametrization Φ_i for a framed 0-sphere along a link, which send the center point of the band transversely across an arc in $\mathcal{A}$ (Move (14)). We finally remark that throughout the process, we may need to change the gradient like vector fields used, but for a fixed Morse function f and choice of regular values $\mathbf{b}$, according to Lemma 10.24, any two gradient like vector fields, v and v' yield parametrized Kirby decompositions which are weakly equivalent.

We make an additional remark about Move (14), where a critical point of Σ crosses $\mathcal{A}$. By a sequence of Moves (10) and (11), we can isotope the Morse function in a neighborhood of the singularity to arrange for it to take the form shown in Figure 10.1. For such a path of Morse functions f_t on Σ yielding the singularity shown in Figure 10.1, the surfaces involved in the new elementary parametrized cobordisms $\mathcal{W}_i'$ and $\mathcal{W}_{i+1}'$, appearing in the statement of Move (14), are the ones labeled in that Figure. See also Figure 11.4.

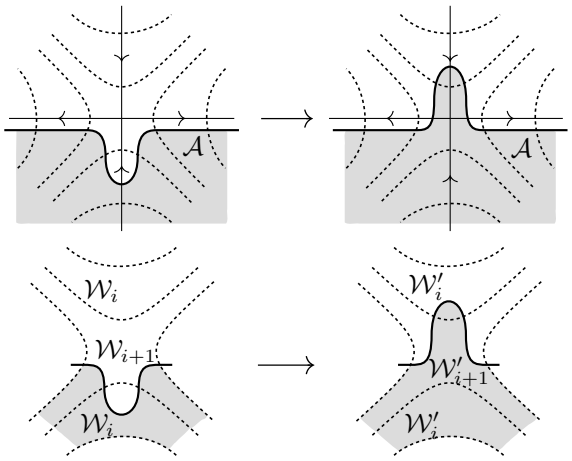

FIGURE 10.1. A local model for Move (14). By possibly composing the path of Morse functions with an isotopy of the surface, we can assume that a singularity of type (5) takes the local model shown on top. The thick solid lines indicate $\mathcal{A}$, while the dashed lines indicate the level sets of f_t, and the thin solid lines with arrows are upward gradient flowlines of f_t. On the bottom, we indicate the change in associated parametrized Kirby decomposition described in the definition of the move. The dotted lines correspond to level sets of the Morse function.

□

11. Construction and invariance of the cobordism maps

In this section, we provide a precise construction of the maps $F_{W,F,\mathfrak{s}}$, for a link cobordism (W, F), and prove invariance from choices involved in the construction.

To define maps associated to a link cobordism (W, F), we take a parametrized Kirby decomposition, $\mathcal{K}$, of (W, F). If we can define maps for each elementary parametrized link cobordism, then we can simply define $F_{W,F,\mathfrak{s}}$ to be the composition of all of them. Note that since only one term of a parametrized Kirby decomposition is allowed to have 2-handles, this agrees with the composition law. In the following subsections, we will combine the maps defined in previous sections for quasi-stabilizations, framed spheres and bands to define maps associated to elementary parametrized Kirby decompositions. Then in Subsection 11.4, we prove invariance from the choice of parametrized Kirby decomposition.

11.1. **Maps induced by elementary parametrized link cobordisms of type** (1) **and** (2)**.** We now consider elementary parametrized link cobordisms $(W, F, \Phi, \mathbb{S}) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ where $\mathbb{S}$ is a framed 0-sphere, 2-sphere or 1-dimensional link which doesn't intersect L (corresponding to type (2) cobordisms), or the empty set (corresponding to type (1) cobordisms). In both cases, the set of divides is a collection of arcs from L_1 to L_2 and the underlying surface of F is a disjoint union of annuli.

Let $\mathcal{A}_0$ denote the pull back of the divides $\mathcal{A}$ to $L_1 \times [0,1]$, under the diffeomorphism Φ. Now $\mathcal{A}_0$ is clearly isotopic relative $L_1 \times \{0,1\}$ to a set of divides, $\mathcal{A}_0'$, which are transverse to each $L_1 \times \{t\}$. The divides $\mathcal{A}_0'$ uniquely specify an isotopy of the basepoints $\mathbf{w}_1$ and $\mathbf{z}_1$, along L_1. We then let $\psi_t : (Y_1, L_1) \to (Y_1, L_1)$ denote an extension of this isotopy to all of Y_1. Setting $\psi = \psi_1$, we define the map

$$F_{W,F,\Phi,\mathbb{S},\mathfrak{s}} = (\Phi|_{Y_1(\mathbb{S})})_* F_{Y_1,\mathbb{S},\Phi^{-1}(\mathfrak{s})} \psi_*.$$

We need to show this is invariant under the choice of isotopic $\mathcal{A}_0'$, as well as invariant under isotopies of Φ, relative Y_1. Although straightforward in our current setting, we prove the following, since we will have to prove a less obvious analog for surfaces with bands later:

Definition 11.1. *If $\mathcal{A}_t$ is a 1-parameter family of dividing sets on $L \times [0,1]$, which are fixed on $L \times \{0,1\}$, we say that $\mathcal{A}_t$ is an **admissible isotopy** if every arc going from $L \times \{0\}$ to $L \times \{1\}$ is transverse to the sets $L \times \{s\}$ for all s.*

Lemma 11.2. *Suppose $(L \times [0,1], \mathcal{A})$ is a cylinder with dividing set, which is a trivial cylinder (i.e. all arcs go from $L \times \{0\}$ to $L \times \{1\}$) or is of type I^+, I^-, II^+ or II^- (there is a single arc which doesn't go from $L \times \{0\}$ to $L \times \{1\}$). Suppose $\mathcal{A}_1$ and $\mathcal{A}_2$ are isotopic to $\mathcal{A}$ through isotopies fixing $L \times \{0,1\}$, and suppose that all of the arcs of $\mathcal{A}_1$ and $\mathcal{A}_2$ which go from $L \times \{0\}$ to $L \times \{1\}$ are transverse to the sets $L \times \{s\}$ for all s. Then $\mathcal{A}_1$ and $\mathcal{A}_2$ are isotopic relative $L \times \{0,1\}$ through a path of dividing sets $\mathcal{A}_t$ such that for all t, all arcs of $\mathcal{A}_t$ which go from $L \times \{0\}$ to $L \times \{1\}$ are transverse to the level sets $L \times \{s\}$.*

Proof. For notational simplicity, assume that L has only one component; the argument easily adapts to the case of multiple components. Suppose $A \in \mathcal{A}$ is any of the arcs going from $L \times \{0\}$ to $L \times \{1\}$ (the set of such arcs is nonempty, since by assumption the links on each ends have basepoints on each end). Given two dividing sets $\mathcal{A}_1$ and $\mathcal{A}_2$ on $L \times [0,1]$, which are both of type I^+, I^-, II^+, II^-, and have the same endpoints, and such that two endpoints of arcs are connected by an arc in $\mathcal{A}_1$ iff they are connected by an arc in $\mathcal{A}_2$, then it is easy to that they are admissibly isotopic iff they share the same algebraic intersection number $\#(A \cap (\{p\} \times [0,1]))$ for any $p \in L$ which is not the endpoint of an arc on either end. The algebraic intersection number is unchanged by arbitrary isotopies which fix $L \times \{0,1\}$, and hence two dividing sets are isotopic relative $L \times \{0,1\}$ iff they are admissibly isotopic relative $L \times \{0,1\}$. □

Lemma 11.3. *Isotopies of the parameterizing diffeomorphism Φ through diffeomorphisms fixing $Y_1 \times \{0\}$ do not affect $F_{W,F,\Phi,\mathbb{S},\mathfrak{s}}$ for type* (1) *and* (2) *elementary cobordisms.*

Proof. Isotopies of Φ which fix Σ have no effect. Similarly isotopies of Φ which are supported only in a neighborhood of Σ, but which fix $L_1 \times \{1\}$ only affect the composition by changing the pullback of the dividing set, and hence possibly changing the map ψ. By Lemma 11.2, such an isotopy of the dividing set can be taken to be an admissible isotopy, which clearly has no effect (since the resulting ψ_* is unchanged). It remains to just check the effect of isotopies supported in a neighborhood of $L_1 \times \{1\}$. If τ denotes a diffeomorphism which twists a neighborhood of $L_1 \times \{1\}$, then precomposing Φ with τ changes $(\Phi|_{Y_1 \times \{1\}})_*$ to $(\Phi|_{Y_1 \times \{1\}})_*(\tau_*)$ but also changes ψ_* to $(\tau^{-1})_*(\psi_*)$ (since the pullback under Φ of the dividing set changes). Since $(\tau^{-1})_*$ commutes with $F_{Y_1,\mathbb{S},\Phi^{-1}(\mathfrak{s})}$, as τ doesn't move the framed sphere $\mathbb{S}$, the factors of $(\tau)_*$ and $(\tau)_*^{-1}$ cancel and the composition is unchanged. □

11.2. **Maps induced by elementary parametrized link cobordisms of type** (3)**.** Suppose $(W, F, \Phi, \mathbb{S}) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a parametrized link cobordism of type (3). Define $\mathcal{A}_0 = \Phi^{-1}(\mathcal{A})$, a collection of divides on $L_1 \times [0,1]$.

Recall that there are four subtypes of elementary cobordisms of type (3), which we denote by I^+, I^-, II^+, and II^-. These are shown in Figure 11.1.

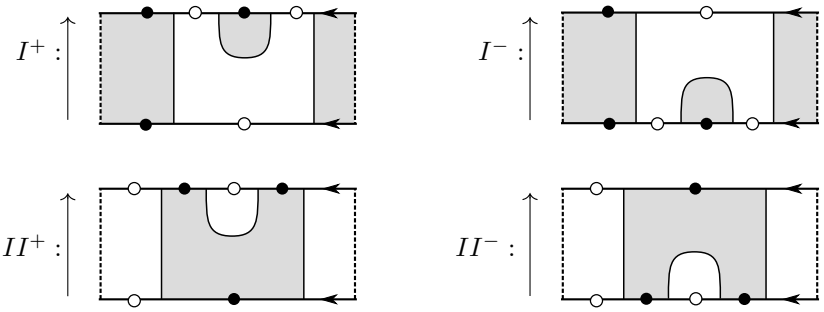

FIGURE 11.1. The four subtypes of elementary cobordisms of type (3).

11.2.1. *Type I^+ elementary cobordisms.* Suppose $(W, F, \Phi, \varnothing) : (Y_1, (L_1, \mathbf{w}_1, \mathbf{z}_1)) \to (Y_2, (L_2, \mathbf{w}_2, \mathbf{z}_2))$ is a parametrized cylindrical cobordism of type $(3I^+)$.

There is a $\mathbf{z}$-region which has three basepoints on its boundary. Let $z_0 \in L_1$ be the basepoint on the incoming end, and let z and z' be the two basepoints on the outgoing end. Let w be the basepoint between z and z', and suppose the basepoints appear with in the following order: z, w then z'. One then isotopes the divides $\mathcal{A}_0$, relative $L_1 \times \{0,1\}$ to get isotopic divides $\mathcal{A}_0'$ such that all arcs from $L_1 \times \{0\}$ to $L_1 \times \{1\}$ are transverse to the sets $L_1 \times \{t\}$. The divides $\mathcal{A}_0'$ determine two diffeomorphisms, ψ and ψ', of Y, which are induced by twisting the link L_1 according to the divides $\mathcal{A}_0'$. The diffeomorphisms ψ and ψ' satisfy

$$\psi(z_0) = z', \qquad \text{and} \qquad \psi'(z_0) = z.$$

By Lemma 11.2, the diffeomorphisms ψ and ψ' are independent (up to isotopies fixing the basepoints) of the choice of $\mathcal{A}_0'$ which is isotopic to $\mathcal{A}_0$.

We now define two maps. The first is

$$F_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* S^+_{w,z} \psi_*$$

and the second is

$$F'_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1 \times \{1\}})_* S^+_{z',w} \psi'_*.$$

Note that there really is no canonical reason to choose one map over the other, since we have to choose which of z and z' we want to identify z_0 with. However an application of Lemma 3.23 shows that the two maps are filtered chain homotopic, so we denote the common map by $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$.

11.2.2. *Type I^- elementary cobordisms.* Suppose $(W, F, \Phi, \varnothing)$ is of type I^-. As with elementary cylindrical cobordisms of type I^-, the divides $\mathcal{A}_0 = \Phi^{-1}(\mathcal{A})$ determine two well defined diffeomorphisms ψ and ψ' of Y, each coming from one of the two maximal identification of a subset of the basepoints in $\mathbf{w}_1 \cup \mathbf{z}_1$ with $\Phi^{-1}(\mathbf{w}_2 \cup \mathbf{z}_2)$ in such a way which respects the divides. Suppose that z and z' are the basepoints on $(L_1, \mathbf{w}_1, \mathbf{z}_1)$ which are in the same region of $\Sigma \setminus \mathcal{A}$, and let w be the basepoint between them. There is one basepoint, $z_0 \in \Phi^{-1}(\mathbf{w}_2 \cup \mathbf{z}_2)$, in the same region on as z and z' on $(L_1 \times [0,1]) \setminus \mathcal{A}_0$, and assume that the ordering is z, w then z'. Assume that

$$\psi(z') = z_0 \qquad \text{and} \qquad \psi'(z) = z_0.$$

We then define two maps:

$$F_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1\times\{1\}})_* S^-_{\psi(w),\psi(z)}\psi_*, \qquad \text{and} \qquad F'_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1\times\{1\}})_* S^-_{\psi'(z'),\psi'(w)}\psi'_*.$$

By Lemma 3.24, these two maps are filtered chain homotopic, so we denote the common map by $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$.

11.2.3. *Type II^+ elementary cobordisms.* Suppose $(W,F,\Phi,\varnothing)$ is of type II^+. Let $\mathcal{A}_0 = \Phi^{-1}(\mathcal{A})$. In the region of $(L_1\times[0,1])\setminus\mathcal{A}_0$ which has three $\mathbf{w}$-basepoints, let w_0 be the basepoint on $L_1\times\{0\}$, and let w and w' be the basepoints on $L_1\times\{1\}$, in this region. Let z denote the basepoint between w and w'. Assume that on $L_1\times\{0\}$, we have the following ordering of basepoints: w', z then w. Using Lemma 11.2, the divides $\mathcal{A}_0$ determine two diffeomorphisms, ψ and ψ', of Y corresponding to twisting a link component as prescribed by the divides. Assume that

$$\psi(w_0) = w' \qquad \text{and} \qquad \psi'(w_0) = w.$$

We define two maps:

$$F_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1\times\{1\}})_* T^+_{w,z}\psi_*, \qquad \text{and} \qquad F'_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1\times\{1\}})_* T^+_{z,w'}\psi'_*.$$

By Lemma 3.25, these two maps are filtered chain homotopic, so we denote the common map by $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$.

11.2.4. *Type II^- elementary cobordisms.* Finally, suppose $(W,F,\Phi,\varnothing)$ is of type II^-. Let $\mathcal{A}_0 = \Phi^{-1}(\mathcal{A})$. In the region of $(L_1\times[0,1])\setminus\mathcal{A}_0$ which has three $\mathbf{w}$-basepoints, let w_0 be the basepoint on $L_1\times\{1\}$, and let w and w' be the basepoints on $L_1\times\{0\}$, in this region. Let z denote the basepoint between w and w', and suppose that the basepoints appear in the order w', z then w. Using Lemma 11.2, the divides $\mathcal{A}_0$ determine two diffeomorphisms, ψ and ψ', of Y corresponding to twisting a link component as prescribed by the divides. Assume that

$$\psi(w') = w_0 \qquad \text{and} \qquad \psi'(w) = w_0.$$

We define two maps:

$$F_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1\times\{1\}})_* T^-_{\psi(w),\psi(z)}\psi_* \qquad \text{and} \qquad F'_{W,F,\Phi,\varnothing,\mathfrak{s}} = (\Phi|_{Y_1\times\{1\}})_* T^-_{\psi'(z),\psi'(w')}\psi'_*.$$

By Lemma 3.26, these two maps are filtered chain homotopic, so we denote the common map by $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$.

Since the maps $S^\pm_{w,z}$ and $T^\pm_{w,z}$ commute with maps induced by diffeomorphisms, adapting Lemma 11.3 immediately yields the following:

Lemma 11.4. *Isotopies of the parameterizing diffeomorphism Φ through diffeomorphisms fixing $Y_1\times\{0\}$, do not affect the maps $F_{W,F,\Phi,\varnothing,\mathfrak{s}}$ for elementary parametrized cobordisms of type* (3)*.*

Proof. This follows identically to Lemma 11.3, using Lemma 11.2. □

11.3. **Maps induced by elementary parametrized link cobordisms of type** (4)**.** Suppose that $(W,F,\Phi,\mathbb{S})$ is an elementary cobordism of type (4) from $(Y_1,\mathbb{L}_1)$ to $(Y_2,\mathbb{L}_2)$. Here

$$\Phi : W(Y_1,L_1,\mathbb{S}) \to (W,F)$$

is the parameterizing diffeomorphism, such that $\Phi^{-1}(\mathcal{A})$ doesn't intersect B, and (it could however, wind around many times on the subsets of $\Sigma(\mathbb{S})$ identified with $L_1\times[0,1]$ or $L_1(\mathbb{S})\times[1,2]$) and the set of divides $\mathcal{A}$ consists entirely of arcs going from the incoming boundary to the outgoing boundary.

Recall the model surface $\Sigma(\mathbb{S}) = (L_1\times[0,1])\cup B\cup(L_1(\mathbb{S})\times[1,2])$ inside of $W(Y_1,\mathbb{S})$ is a diffeomorphic to a disjoint union of a pair of pants together with some cylinders. There is a subset, B, identified with $[-1,1]\times[-1,1]$. Define the set

$$\mathcal{L} = (L\times\{1\})\cap(L(\mathbb{S})\times\{1\}),$$

which can be thought of as the level set, outside of the band, of a Morse function with a single index one critical point on the band.

Definition 11.5. *We will call an isotopy $\mathcal{A}_t$ of divides on $\Sigma(\mathbb{S})$, which is fixed on $L\times\{0\}$ and $L(\mathbb{S})\times\{2\}$, an **admissible isotopy** if each arc of $\mathcal{A}_t$ intersects $\mathcal{L}$ at a exactly one point, transversely, and no arcs ever intersect B.*

Lemma 11.6. *Suppose that $\mathcal{A}_1$ and $\mathcal{A}_2$ are collections of divides on $\Sigma(\mathbb{S}) = (L \times [0,1]) \cup B \cup (L(\mathbb{S}) \times [1,2])$ and that each arc of both divides goes from $L \times \{0\}$ to $L(\mathbb{S}) \times \{2\}$. Suppose further that each arc of $\mathcal{A}_1$ and $\mathcal{A}_2$ intersects $\mathcal{L}$ exactly once, transversely, and that no arcs of either $\mathcal{A}_i$ intersect B. If $\mathcal{A}_1$ and $\mathcal{A}_2$ are isotopic through divides $\mathcal{A}_t$ which do not intersect the point $(0,0) \in B$ and are fixed on $L \times \{0\}$ and $L(\mathbb{S}) \times \{2\}$, then $\mathcal{A}_1$ and $\mathcal{A}_2$ are admissibly isotopic.*

Proof. On any component of $\Sigma(\mathbb{S})$ which doesn't intersect B, the claim is obvious, so suppose without loss of generality that L is a one (resp. two) component link which is separated (resp. joined) by the band B. Suppose for definiteness, that L is a two component link which is joined by B. The case that L is a one component link which is separated by the band follows from nearly identical reasoning.

Let $\{p_-, p_+\} \subseteq L$ denote the underlying 0-sphere of $\mathbb{S}$. Let C_- and C_+ be curves on $(L \times [0,1]) \cup B$, each consisting of a path on $L \times [0,1]$ from a point on $L \times \{0\}$ which is not an endpoint of $\mathcal{A}_i$, to $(p_-, 1)$ or $(p_+, 1)$, respectively, concatenated with one of the arcs $[-1,0] \times \{0\}$ or $[0,1] \times \{0\}$, respectively, on B. These are shown in Figure 11.2.

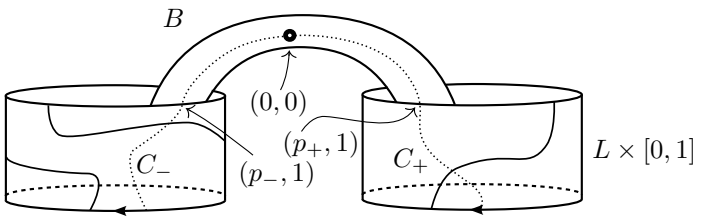

FIGURE 11.2. The arcs C_- and C_+ on $\Sigma(\mathbb{S})$, as well as the band B and the point $(0,0) \in B$.

Similarly let C'_+ and C'_- be two similarly defined curves on $B \cup (L(\mathbb{S}) \times [1,2])$ from $(0,0)$ to two points on $L(\mathbb{S}) \times \{2\}$ which are not endpoints of any of the arcs of $\mathcal{A}_i$.

The admissible isotopy class of the divides $\mathcal{A}_1$ and $\mathcal{A}_2$ relative $L \times \{0\} \cup L(\mathbb{S}) \times \{2\}$ is determined by the ordering that the arcs appear on $\mathcal{L}$, and the intersection numbers of arcs with C_-, C_+, C'_-, and C'_+.

The configuration of endpoints of arcs on $L \times \{1\}$ is determined by the region of $\Sigma \setminus \mathcal{A}$ that the band is in and the order that the arcs appear on $L \times \{0\}$. This is unchanged by isotopies of $\Sigma(\mathbb{S})$ which may not be admissible, but which don't map any arcs to $(0,0)$. Similarly the algebraic intersection numbers mentioned above are not changed by isotopies of the arcs which fix the ends of the arcs on $L \times \{0\}$ and $L(\mathbb{S}) \times \{2\}$, and don't map any arcs to $(0,0)$. Hence if two collections of divides, $\mathcal{A}_1$ and $\mathcal{A}_2$, each consisting of arcs going from $L \times \{0\}$ to $L(\mathbb{S}) \times \{2\}$ and sharing the same end points, are isotopic relative $L \times \{0\}$ and $L(\mathbb{S}) \times \{2\}$ through isotopies not sending an arc to $(0,0)$, then they are admissibly isotopic.

□

We now arbitrarily pick basepoints $\mathbf{w}_{1.5} \cup \mathbf{z}_{1.5}$ on the $L_1 \times \{1\}$ which make $(L_1 \times \{1\}, \mathbf{w}_{1.5}, \mathbf{z}_{1.5})$ an acceptably based link, and which make the band B either an α-band or a β-band. Note that the type of the band as either type-$\mathbf{w}$ or type-$\mathbf{z}$ is well defined by the configuration of the divides. For a choice of $\mathbf{w}_{1.5}$ and $\mathbf{z}_{1.5}$ there is no ambiguity in whether B is an α-band or a β-band, but up to isotopy which fixes B, there are two choices of basepoints $\mathbf{w}_{1.5}$ and $\mathbf{z}_{1.5}$ which respect the divides (these two choices are isotopic through isotopies which don't fix B). Let $\mathbf{w}^\tau_{1.5}$ and $\mathbf{z}^\tau_{1.5}$ denote these basepoints, for a choice of $\tau \in \{\alpha, \beta\}$.

Adapting Lemma 11.2 in the obvious way, the divides $\Phi^{-1}(\mathcal{A})$ on $L_1 \times [0,1]$ induce an isotopy ψ^τ_1 of (Y, L) which sends $(\mathbf{w}_1, \mathbf{z}_1)$ to $(\mathbf{w}^\tau_{1.5}, \mathbf{z}^\tau_{1.5})$. Similarly the divides $\Phi^{-1}(\mathcal{A})$ induce an isotopy ψ^τ_2 which sends $(\mathbf{w}^\tau_{1.5}, \mathbf{z}^\tau_{1.5})$ to $(\Phi^{-1}(\mathbf{w}_2), \Phi^{-1}(\mathbf{z}_2))$. We now define, with respect to these choices

$$F^\tau_{W,F,\Phi,\mathbb{S},\mathfrak{s}} = \Phi|_{Y_1(\mathbb{S}) \times \{2\}}(\psi^\tau_2)_* F^{\mathbf{o}}_{Y_1,\mathbb{L},\mathbb{S},\Phi^{-1}(\mathfrak{s})}(\psi^\tau_1)_*.$$

There is no ambiguity in the choice of $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$, though τ can be either α or β, as determined by the choice of basepoints $\mathbf{w}^\tau_{1.5}$ and $\mathbf{z}^\tau_{1.5}$. By Propositions 7.8 and 7.11 we know that the α- and β-band maps are related by the diffeomorphism resulting from moving basepoints, which is tautologically encoded into ψ^τ_1 and ψ^τ_2. From these considerations we conclude that

$$F^\alpha_{W,F,\Phi,\mathbb{S},\mathfrak{s}} \simeq F^\beta_{W,F,\Phi,\mathbb{S},\mathfrak{s}},$$

and we denote the common map

$$F_{W,F,\Phi,\mathbb{S},\mathfrak{s}}.$$

We have the following:

Lemma 11.7. *For an elementary parametrized cobordism $(W,F,\Phi,\mathbb{S})$ of type (4), isotopies of the parametrization Φ which fix $Y_1\times\{0\}$ and don't map the point $(0,0)\in B$ to $\mathcal{A}$ do not affect the map $F_{W,F,\Phi,\mathbb{S},\mathfrak{s}}$.*

Proof. Isotopies which are fixed on $\Sigma(\mathbb{S})$ clearly have no effect, so we can just consider isotopies of $\Phi|_{\Sigma(\mathbb{S})}$, extended to isotopies of all of W. By the same argument as in Lemma 11.3, the maps are independent of isotopies of Φ which are supported in a neighborhood of $L_1(\mathbb{S})\times\{2\}$. Hence it remains to see that the maps are invariant under isotopies which are fixed on $Y_1\times\{0\}$ and $Y_1(\mathbb{S})\times\{2\}$. Since Φ_t is fixed on $Y(\mathbb{S})$, the map $F_{W,F,\Phi,\mathbb{S},\mathfrak{s}}$ only depends on Φ through the dividing set $\Phi^{-1}(\mathcal{A})$. If Φ_t is such an isotopy, we consider the dividing sets $\mathcal{A}_t=\Phi_t^{-1}(\mathcal{A})$ on $\Sigma(\mathbb{S})$. By Lemma 11.6, we can assume that $\mathcal{A}_0$ and $\mathcal{A}_1$ are admissibly isotopic, i.e. the dividing sets $\mathcal{A}_t$ never pass through the band B, and the curves $\mathcal{A}_t$ remain transverse to $\mathcal{L}$ throughout.

Such an isotopy of $\mathcal{A}_t$ induces an isotopy of both (Y_1,L_1) and $(Y_1(\mathbb{S}),L_1(\mathbb{S}))$ which fixes $\mathbb{S}\subseteq Y$ and $D^1\times S^2\subseteq Y(\mathbb{S})$, respectively, resulting from moving the basepoints according to the intersection of $\mathcal{A}_t$ with $\mathcal{L}$. We let τ denote this isotopy on Y_1, and let $\tau^{\mathbb{S}_1}$ denote this isotopy on $Y_1(\mathbb{S})$. The effect of the isotopy of $\mathcal{A}$ is to replace ψ_1 with $\psi_1'=\tau\circ\psi_1$ and to replace ψ_2 with $\psi_2'=\psi_2\circ(\tau^{\mathbb{S}_i})^{-1}$. In general, the band maps $F_B^{\mathbf{o}}$ are not invariant under diffeomorphisms which don't fix the band. Instead we canonically have $\phi_*F_B=F_{\phi_*B}\phi_*$ for a diffeomorphism ϕ. Since τ fixes $\mathbb{S}\subseteq Y$, and $\tau^{\mathbb{S}}$ fixes $D^1\times S^2\subseteq Y(\mathbb{S})$, we have in particular that $\tau^{\mathbb{S}}$ fixes the band $B\subseteq Y_1(\mathbb{S})$. Hence we have

$$\begin{aligned}\Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi_2')_*F^{\mathbf{o}}_{Y_1,\mathbb{L},\mathbb{S},\mathfrak{s}}(\psi_1')_*&\simeq\Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi_2)_*(\tau^{\mathbb{S}_i})_*^{-1}F^{\mathbf{o}}_{Y,\mathbb{L},\mathbb{S},\mathfrak{s}}\tau_*(\psi_1)_*\\ \simeq\Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi_2)_*F^{\mathbf{o}}_{Y_1,\mathbb{L},\mathbb{S},\mathfrak{s}}(\tau)_*^{-1}\tau_*(\psi_1)_*&\simeq\Phi|_{Y_1(\mathbb{S})\times\{2\}}(\psi_2)_*F^{\mathbf{o}}_{Y_1,\mathbb{L},\mathbb{S},\mathfrak{s}}(\psi_1)_*,\end{aligned}$$

showing that the maps are unchanged. □

11.4. **Definition and invariance of the cobordism maps.** If (W,F) is a link cobordism such that each component of W has a nonempty incoming and outgoing end, and each component of Σ has a nonempty incoming and outgoing end, we take a parametrized Kirby decomposition, $\mathcal{K}$, with elementary parametrized cobordisms $\mathcal{W}_i=(W_i,F_i,\Phi_i,\mathbb{S}_i)$, and define the maps

$$F_{W,F,\mathfrak{s},\mathcal{K}}=\prod_i F_{W_i,F_i,\Phi_i,\mathbb{S}_i,\mathfrak{s}|_{W_i}},$$

where the maps $F_{W_i,F_i,\Phi_i,\mathbb{S}_i,\mathfrak{s}|_{W_i}}$ are the maps defined in the previous subsections.

We now prove a slightly weaker version of Theorem A, for cobordisms which have nonempty incoming and outgoing ends of each component of W and Σ. We will extend this to the most general version of Theorem A in Section 12.

Theorem A′. *The maps $F_{W,F,\mathfrak{s},\mathcal{K}}$ are independent of the parametrized Kirby decomposition, $\mathcal{K}$, up to filtered, equivariant chain homotopy, and hence define invariants of the decorated link cobordisms with* Spin^c *structure, $(W,F,\mathfrak{s})$.*

Proof. By Proposition 10.3, it is sufficient to show that the maps $F_{W,F,\mathfrak{s},\mathcal{K}}$ are equal for weakly equivalent parametrized Kirby decompositions. Hence it is sufficient to check invariance from the moves of Definition 10.2.

Move 1: We consider strong isotopies (Definition 10.1). First note that if

$$\phi:(W,\Sigma)\to(W,\Sigma)$$

is any diffeomorphism with $\phi|_{\partial W}=\mathrm{id}$, and $\mathcal{K}$ is a parametrized Kirby decomposition for (W,F), then using Equation (10) we get a decomposition $\phi_*\mathcal{K}$ of (W,F) into parametrized cobordisms, though in general, as we remarked after Definition 10.1, the decomposition may fail to be a parametrized Kirby decomposition for (W,F). Instead, tautologically, $\phi_*\mathcal{K}$ yields a parametrized Kirby decomposition for (W,ϕ_*F), where $\phi_*F=(\Sigma,\phi_*\mathcal{A})$. Indeed, we tautologically have that

$$F_{W,F,\mathfrak{s},\mathcal{K}}\simeq F_{W,\phi_*F,\phi_*\mathfrak{s},\phi_*\mathcal{K}},$$

or more explicitly

$$F_{W,(\Sigma,\mathcal{A}),\mathfrak{s},\mathcal{K}} \simeq F_{W,(\Sigma,\phi_*\mathcal{A}),\phi_*\mathfrak{s},\phi_*\mathcal{K}}.$$

In the case where $\phi = \phi_1$ for a strong isotopy ϕ_t, we note that our desired equality now becomes

$$F_{W,(\Sigma,(\phi_1)_*\mathcal{A}),(\phi_1)_*\mathfrak{s},(\phi_1)_*\mathcal{K}} \simeq F_{W,(\Sigma,\mathcal{A},)(\phi_1)_*\mathfrak{s},(\phi_1)_*\mathcal{K}}.$$

Hence, to show invariance under strong isotopies, it is sufficient to show that isotopies of the arcs in $\mathcal{A}$, which are fixed on ∂W, which additionally preserve the property that each term in the decomposition is an elementary parametrized cobordism, do not affect the maps. Such an isotopy can be arranged into a sequence of isotopies which are supported in a neighborhood of an individual Y_i, and isotopies which are fixed on all of the Y_i.

For isotopies of $\mathcal{A}$ which are fixed on all of the Y_i, invariance follows from Lemmas 11.3, 11.4, and 11.7. Hence it remains only to show invariance under isotopies of $\mathcal{A}$ which are supported in a neighborhood of a fixed Y_i which is in the interior of W. For such isotopies, let $\mathcal{W}_{i-1} = (W_{i-1}, F_{i-1}, \Phi_{i-1}, \mathbb{S}_{i-1})$ and $\mathcal{W}_i = (W_i, F_i, \Phi_i, \mathbb{S}_i)$ denote the elementary parametrized cobordisms adjacent to Y_i for $\mathcal{K}$. Let $\tau : W \to W$ be the result of taking an isotopy τ_t of L_i, and extending it to an isotopy of Y_i, then extending that to an isotopy of all of W which preserves Σ, and setting $\tau = \tau_1$. It is straightforward to check that for each type of elementary parametrized cobordism that

$$F_{W_{i-1},(\Sigma_{i-1},\tau(\mathcal{A})\cap\Sigma_{i-1}),\Phi_{i-1},\mathfrak{s}|_{W_{i-1}}} \simeq \tau_* \circ F_{W_{i-1},(\Sigma_{i-1},\mathcal{A}\cap\Sigma_{i-1}),\Phi_{i-1},\mathfrak{s}|_{W_{i-1}}}$$

and

$$F_{W_i,(\Sigma_i,\tau(\mathcal{A})\cap\Sigma_i),\Phi_i,\mathfrak{s}_i} \simeq F_{W_i,(\Sigma_i,\mathcal{A}_i),\Phi_i,\mathfrak{s}_i} \circ (\tau^{-1})_*.$$

When the two maps are composed, the effects of τ cancel and the composition is unchanged.

Move 2: We now consider isotopies of a framed sphere, away from L_i. We will omit the subscript i, and the Spin^c structures, for notational simplicity.

Suppose that $\mathbb{S}^t$ is an isotopy of framed spheres away from L, starting at $\mathbb{S}$ and ending at $\mathbb{S}'$. Let d be the (isotopically unique) diffeomorphism of Y extending the canonical diffeomorphism sending $\mathbb{S}$ to $\mathbb{S}'$ along this isotopy. Let $d^{\mathbb{S}} : Y(\mathbb{S}) \to Y(\mathbb{S}')$ denote the induced map. Let ψ denote the diffeomorphism of (Y, L) obtained by pulling back the divides under Φ to the model cobordism $W(Y, L, \mathbb{S})$, and defining ψ to be the map induced by twisting along L, according to the divides. Notice that Φ and Φ' can be taken to be equal along $L \times [0, 1] \subseteq W(Y, L, \mathbb{S})$, since the isotopy of framed spheres occurs away from L. By definition we have that

$$F_{W,\mathbb{L},\mathbb{S},\Phi} = (\Phi|_{Y(\mathbb{S})})_* F_{Y,\mathbb{L},\mathbb{S}}\psi_*$$

and

$$F_{W,\mathbb{L},\mathbb{S}',\Phi'} = (\Phi'|_{Y(\mathbb{S}')})_* F_{Y,\mathbb{L},\mathbb{S}'}\psi_*.$$

By definition of the move, we have that

$$(\Phi'|_{Y(\mathbb{S}')})_* d^{\mathbb{S}}_* = (\Phi|_{Y(\mathbb{S})})_*.$$

Similarly, by tautological we have

$$F_{Y,\mathbb{L},\mathbb{S}'} d_* = d^{\mathbb{S}}_* F_{Y,\mathbb{L},\mathbb{S}}.$$

We note that ψ_* and d_* commute, since they are induced by isotopies supported on disjoint subsets of Y. As a consequence we have that

$$\begin{aligned} F_{W,\mathbb{L},\mathbb{S}',\Phi'} d_* &= (\Phi'|_{Y(\mathbb{S}')})_* F_{Y,\mathbb{L},\mathbb{S}'}\psi_* d_* \\ &= (\Phi'|_{Y(\mathbb{S}')})_* F_{Y,\mathbb{L},\mathbb{S}'} d_* \psi_* \\ &= (\Phi'|_{Y(\mathbb{S}')})_* d^{\mathbb{S}}_* F_{Y,\mathbb{L},\mathbb{S}}\psi_* \\ &= (\Phi|_{Y(\mathbb{S})})_* F_{Y,\mathbb{L},\mathbb{S}}\psi_* \\ &= F_{W,\mathbb{L},\mathbb{S},\Phi}. \end{aligned}$$

Now d is a diffeomorphism of Y which is isotopic to the identity relative L, and hence we conclude that $d_* \simeq 1$. Hence

$$F_{W,\mathbb{L},\mathbb{S}',\Phi'} \simeq F_{W,\mathbb{L},\mathbb{S},\Phi},$$

showing the maps to be invariant under this move.

Move 3: We now prove invariance of isotopies of a framed 0-sphere along a link L_i. This follows similarly to Move (2), and is essentially a tautology. Suppose that $(W, F, \mathbb{S}, \Phi)$ is an elementary parametrized link cobordism of type (4). The type, $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$, of the elementary parametrized cobordism is determined uniquely by the type of the region of F containing the image of $(0,0) \in B$, but by construction the image of the point $(0,0)$ in the model band does not change. For notational simplicity we omit the i subscript. Let d_t be an isotopy of $\mathbb{S}$ along L, define $d = d_1$ and let $\mathbb{S}'$ be the image of $\mathbb{S}$ and let Φ' be the parametrization described in Move (3). Let ψ_1 and ψ_2 be the diffeomorphisms in the definition of the maps (cf. Subsection 11.3) associated to Φ, and let ψ_1' and ψ_2' be the diffeomorphisms associated to the parametrization Φ'. By definition

$$F_{W,F,\mathbb{S},\Phi} = (\Phi|_{Y(\mathbb{S})\times\{2\}})_*(\psi_2)_* F^{\mathbf{o}}_{Y,\psi_1(\mathbb{L}),\mathbb{S}}(\psi_1)_*$$

and

$$F_{W,F,\mathbb{S}',\Phi'} = (\Phi'|_{Y(\mathbb{S})\times\{2\}})_*(\psi_2')_* F^{\mathbf{o}}_{Y,\psi_1'(\mathbb{L}),\mathbb{S}'}(\psi_1')_*.$$

It is easy to see that by construction we have that

$$(\psi_1')_* = d_*(\psi_1)_*.$$

Tautologically we have

$$(d^{\mathbb{S}})_*^{-1} F_{Y,\psi_1'(\mathbb{L}),\mathbb{S}'} d_* = F_{Y,\psi_1(\mathbb{L}),\mathbb{S}}.$$

Now also by construction, we have

$$\psi_2' = (d^{\mathbb{S}})_*(\psi_2)_*(d^{\mathbb{S}})_*^{-1}.$$

Finally we note that, by construction

$$(\Phi|_{Y(\mathbb{S})\times\{2\}})_* = (\Phi'|_{Y(\mathbb{S}')\times\{2\}})_* d^{\mathbb{S}}_*.$$

As a consequence, we see that

$$F_{W,F,\mathbb{S}',\Phi'} = (\Phi'|_{Y(\mathbb{S})\times\{2\}})_*(\psi_2')_* F^{\mathbf{o}}_{Y,\psi_1'(\mathbb{L}),\mathbb{S}'}(\psi_1')_* = (\Phi'|_{Y(\mathbb{S})\times\{2\}})_*(d^{\mathbb{S}})_*(\psi_2)_*(d^{\mathbb{S}})_*^{-1} F^{\mathbf{o}}_{Y,\psi_1'(\mathbb{L}),\mathbb{S}'} d_*(\psi_1)_*$$
$$= (\Phi|_{Y(\mathbb{S})\times\{2\}})_*(\psi_2)_* F^{\mathbf{o}}_{Y,\psi_1(\mathbb{L}),\mathbb{S}}(\psi_1)_* = F_{W,F,\mathbb{S},\Phi},$$

completing the proof of invariance under Move (3).

Move 4: Invariance under adding or removing a trivial cylinder is essentially a tautology. The map induced by a trivial cylinder is a diffeomorphism map, and all of the other maps involve pre- or post-composition with an analogous diffeomorphism. It is easy to check that combining a trivial cylinder with another elementary cobordism (from the left or the right) does not affect the maps.

Move 5: We now consider invariance under changing the orientation of the framed sphere $\mathbb{S}$. We argue in the case of a framed sphere $\mathbb{S}$, attached away from L. The case that $\mathbb{S}$ is attached along L follows from similar considerations. Let $\sigma : \mathbb{R}^4 \to \mathbb{R}^4$ denote the map $\sigma(x, y, w, z) = (-x, y, w, -z)$. If $\iota : S^k \times D^{3-k} \to Y$ is the embedding of $\mathbb{S}$, let $\overline{\mathbb{S}}$ denote the framed sphere with embedding $\iota \circ \sigma$. Note that σ induces a map $\sigma : Y(\mathbb{S}) \to Y(\overline{\mathbb{S}})$, as well as $\sigma : W(Y, L, \mathbb{S}) \to W(Y, L, \overline{\mathbb{S}})$. Indeed one easily sees that

$$\sigma \circ F_{Y,\mathbb{S},\mathfrak{s}} = F_{Y,\overline{\mathbb{S}},\sigma_*\mathfrak{s}},$$

which follows from the fact that the maps don't depend on a choice of orientation for the framed spheres $\mathbb{S}$. Since the parametrizations Φ and Φ' also differ by composition with σ, the composed maps are unchanged.

Move 6: Invariance under index 1/2 critical point birth-deaths of $f|_{W\setminus\Sigma}$ follows from the same triangle map computation used to prove the analogous result about the closed 3-manifold invariants ([Zem15b, Theorem 9.7]). We remark that by construction, the handles are attached by framed spheres which are in the complement of the link, allowing the same construction to be used. We note that the above count of holomorphic triangles is an adaptation of the original argument by Ozsváth and Szabó in [OS06, Lemma 4.16] to handle our slightly different definition of the 1-handle maps than in the original papers.

Move 7: Invariance under index 2/3 critical point birth-deaths of $f|_{W\setminus\Sigma}$ follows from [Zem15b, Theorem 9.11]. Note that although generically a framed 2-sphere in Y could intersect the link at a collection of points, by our construction, any framed 2-sphere appearing in a parametrized Kirby decomposition does not intersect the link L, which allows us to use the same model computation as in the case of [Zem15b, Theorem 9.11]. Again note that an analogous count of holomorphic triangles was proven by Ozsváth and Szabó in [OS06, Lemma 4.17].

Move 8: Invariance under critical value swaps between two index 1 or two index 3 critical points of $f|_{W\setminus\Sigma}$ follows from the commutation of 1- and 3-handle maps amongst each other, proven in Lemma 4.4.

Move 9: Invariance from handleslides amongst the components of a framed 1-dimensional link follows identically as it did in [OS06, Lemma 4.14]. In summary, a handleslide of components of a framed link is reflected on a Heegaard triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ subordinate to an α-bouquet for the framed link by handlesliding the corresponding $\boldsymbol{\alpha}$-curves across each other. Invariance from such handleslide follows using the associativity relations, and a model computation (performed in [OS06]).

Move 10: Invariance from a birth-death singularity of $f|_{\mathcal{A}}$ follows from the identities $S^-_{w,z}T^+_{w,z} \simeq T^-_{w,z}S^+_{w,z} \simeq \mathrm{id}$ proven in Lemma 3.13. This is demonstrated in Figure 11.3. If instead we switched the orientation of the links with respect to this picture, independence would follow by replacing $S^{\pm}_{w,z}$ with $S^{\pm}_{z,w}$ and $T^{\pm}_{w,z}$ with $T^{\pm}_{z,w}$.

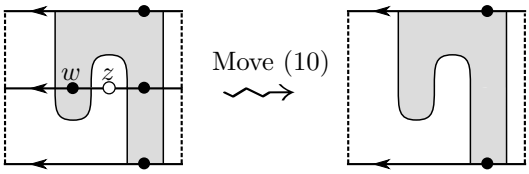

FIGURE 11.3. A picture of Move 10. The map induced on the left is $T^-_{w,z}S^+_{w,z}$ if the cobordism is going from bottom to top, and $S^-_{w,z}T^+_{w,z}$, if the cobordism is going from top to bottom. The map on the right is the identity.

Move 11: For a critical value switch of critical points of $f|_{\mathcal{A}}$, there are 10 possible configurations, depending on whether the two type (3) elementary parametrized cobordisms are of subtype I^+, I^-, II^+, or II^-. The cases where both are of type I^+ or I^- are handled by Theorem 3.15. The cases where both are of type II^+ or II^- are covered by Lemma 3.16, and the cases that one is of type $I^{\pm}$ and the other is of type $II^{\pm}$ is covered by Proposition 3.17.

It is important to note that in this last case, one does not expect $T^{\pm}_{w,z}$ and $S^{\pm}_{w',z'}$ to always commute. Indeed in Proposition 3.17, there were requirements on the configuration of basepoints which can appear. It is easy to check that the requirements on the configuration of basepoints is always satisfied when the basepoints arise in critical value swap of critical points of $f|_{\mathcal{A}}$ where one induces a type $II^{\pm}$ map, and the other induces a type $I^{\pm}$ map (see Figure 11.1 for a picture of type $I^{\pm}$ and $II^{\pm}$ decorations for elementary parametrized cobordisms of type (3)).

Move 12: We now consider the effect of switching the order of two index 1 critical points on Σ. The maps associated to critical points along Σ are a composition of the 1-handle maps, as well as the $\mathbf{z}$-band or $\mathbf{w}$-band maps.

We note that by the triangle map computation of [Zem15b, Theorem 8.8], the 1-handle maps commute with the band maps. By Lemma 4.4 the 1-handle maps commute with each other. Thus it remains just to show that the band maps can be commuted amongst themselves. This now follows from Proposition 7.12. Note that if both band maps are $\mathbf{z}$-band maps or both are $\mathbf{w}$-band maps, then the maps always commute, though if one band map is type-$\mathbf{w}$ and the other is type-$\mathbf{z}$, there are assumptions from Proposition 7.12 about the configuration of ends of the band. These are automatically satisfied in the case that the configuration of bands came from a surface with divides with a Morse function f such that f has two index one critical points, and $f|_{\mathcal{A}}$ has no critical points.

Move 13: We now consider a critical value switch between critical points of $f|_{\Sigma}$ and $f|_{\mathcal{A}}$. This follows from a similar argument as the last move. Again, the 1-handle maps commute with quasi-stabilization maps by Lemma 4.3 and the band maps by the triangle map computation of [Zem15b]Theorem 8.8. Thus it is sufficient to show that the band maps commute with the maps induced by elementary cobordism maps of type (3).

Commutation between the maps $S^{\pm}_{w,z}$ and $T^{\pm}_{w,z}$ and the band maps $F^{\mathbf{w}}_B$ and $F^{\mathbf{z}}_B$ now follows from Propositions 7.1 and 7.3. We note that the hypotheses of those two propositions regarding the configuration of basepoints involved are always satisfied when the configuration of basepoints and bands are induced by a

surface with divides with a Morse function f which has a single index one critical point, and such that $f|_{\mathcal{A}}$ also has a single index 1 critical point.

Move 14: We now consider move corresponding to a critical point of Σ crossing $\mathcal{A}$. Notice that the 1-handle maps commute with everything, so we can bring them below the quasi-stabilization maps and the band maps, so we can just show the equality of maps related to the quasi-stabilization maps and the band maps. In this move, we have arranged for the crossing to locally be as in Figure 10.1. With basepoints, we can redraw this as in Figure 11.4.

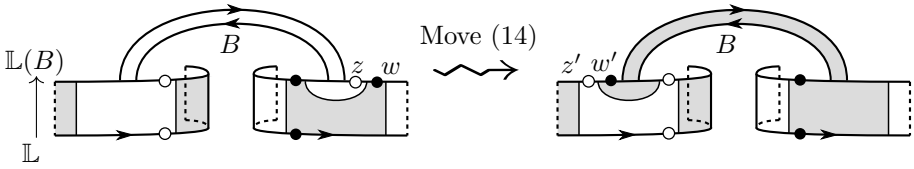

FIGURE 11.4. The two configurations we need to show induce the same map, after moving basepoints, for Move 14.

The maps corresponding to the bands involve postcomposition with a diffeomorphism map which fixes the link and moves the basepoints according to the divides. Hence it is sufficient to show that the indicated maps from Figure 11.4 become equal after moving (w, z) to (w', z'). The map on the right induces

$$F_B^{\mathbf{w}} S^+_{w',z'},$$

while the map on the left induces

$$F_B^{\mathbf{z}} T^+_{w,z}.$$

The diffeomorphism moving (z', w') to (z, w) takes the form $S^-_{w',z'} T^+_{w,z}$ by Lemma 3.29. Thus we wish to show

$$S^-_{w',z'} T^+_{w,z} F_B^{\mathbf{w}} S^+_{w',z'} \simeq F_B^{\mathbf{z}} T^+_{w,z}.$$

This will be an algebraic computation. We note first that

$$\Psi_z T^+_{w,z} \simeq 0, \tag{13}$$

which follows immediately by direct computation using the quasi-stabilized differential in Proposition 3.4 and the definition of $T^+_{w,z}$. We thus compute:

$$\begin{aligned}
S^-_{w',z'} T^+_{w,z} F_B^{\mathbf{w}} S^+_{w',z'} &\simeq S^-_{w',z'} F_B^{\mathbf{w}} T^+_{w,z} S^+_{w',z'} && \text{(Proposition 7.3)}\\
&\simeq S^-_{w',z'} (F_B^{\mathbf{z}} \Psi_z + \Psi_z F_B^{\mathbf{z}}) T^+_{w,z} S^+_{w',z'} && \text{(Proposition 7.15)}\\
&\simeq S^-_{w',z'} \Psi_z F_B^{\mathbf{z}} T^+_{w,z} S^+_{w',z'} && \text{(Equation (13))}\\
&\simeq S^-_{w',z'} \Psi_z F_B^{\mathbf{z}} S^+_{w',z'} T^+_{w,z} && \text{(Proposition 3.17)}\\
&\simeq S^-_{w',z'} \Psi_z S^+_{w',z'} F_B^{\mathbf{z}} T^+_{w,z} && \text{(Proposition 7.1)}\\
&\simeq F_B^{\mathbf{z}} T^+_{w,z}, && \text{(Remark 3.14)}
\end{aligned}$$

completing the proof. □

12. Maps for cobordisms where $\pi_0(\mathbb{L}_i) \to \pi_0(\Sigma)$ is not surjective

In this section we consider cobordisms where $\pi_0(\mathbb{L}_0) \to \pi_0(\Sigma)$ is not surjective. As in [Zem15b], to construct these maps, we will need to pre- or post-compose with 0- and 4-handle maps, which we now define. In [Zem15b], we introduced them in the case that we were adding a copy of S^3 with a single new basepoint. We now wish to define them in the case that we are adding a copy of S^3 containing an unknot with two basepoints.

Definition 12.1. *If $B^4 \subseteq \operatorname{int} W$ is a 4-ball, intersecting Σ along an arc of $\mathcal{A}$, we say B is **standard** if $B \cap \mathcal{A}$ is connected and there are coordinates (x, y, w, z) on a neighborhood of $B \subseteq \operatorname{int} W$ such that $B = \{(x, y, w, z) : x^2 + y^2 + w^2 + z^2 < 1\}$ and in these coordinates $\Sigma \cap B = \{(x, y, 0, 0)\} \cap B$.*

In this section, we will implicitly assume all 4-balls B along $\mathcal{A}$ are standard. Note that as we did in Subsection 8.2, when we constructed the compound 1-handle/band maps, if we pick a vector in the unit sphere bundle $U((T_{(0,0,0,0)}\Sigma)^{\perp}) \cong S^1$, we can push the surface Σ off of B and into ∂B. As S^1 is connected, any two ways of doing this yield isotopic disks in $\partial B = S^3$, relative the unknot $U = \partial B \cap \Sigma$. From another perspective, if $U \subseteq S^3$ is an unknot then any two disks with boundary on U are isotopic relative U, which can be seen from other considerations. With this in mind, we proceed to define the maps for general cobordisms which might not have enough ends.

Let $\mathbb{U} \subseteq S^3$ denote an unknot with exactly two basepoints. We now define the 0-handle and 4-handle maps for link Floer homology:

$$F_{Y,0,\mathbb{U}} : \mathcal{C}FL^{\circ}(Y, \mathbb{L}, \mathfrak{s}) \to \mathcal{C}FL^{\circ}(Y \sqcup S^3, \mathbb{L} \sqcup \mathbb{U}, \mathfrak{s} \sqcup \mathfrak{s}_0)$$

and

$$F_{Y,4,\mathbb{U}} : \mathcal{C}FL^{\circ}(Y \sqcup S^3, \mathbb{L} \sqcup \mathbb{U}, \mathfrak{s} \sqcup \mathfrak{s}_0) \to \mathcal{C}FL^{\circ}(Y, \mathbb{L}, \mathfrak{s}).$$

These take the obvious form, namely

$$F_{Y,0,\mathbb{U}}(\mathbf{x}) = \mathbf{x} \times c \qquad \text{and} \qquad F_{Y,4,\mathbb{U}}(\mathbf{x} \times c) = \mathbf{x}$$

for c the generator of $\mathcal{C}FL^{-}(S^3, \mathbb{U}, \mathfrak{s}_0)$.

Remark 12.2. Earlier, we defined maps for a "compound 1-handle/band", corresponding to attaching a 4-dimensional 1-handle for W, which contained a 2-dimensional 1-handle for Σ. Analogously, the 0-handle and 4-handle maps correspond to attaching a 4-dimensional 0- or 4-handle for W, which contains a 2-dimensional 0- or 2-handle for Σ.

Suppose one attaches a 0-handle to $Y \times [0,1]$, and then attaches a 1-handle with framed 0-sphere $\mathbb{S}_0 = \{p_1, p_2\} \times D^3$, with one foot at a point p_1 in Y and one foot at a point p_2 in the new copy of S^3. There is a diffeomorphism $\phi : (Y \sqcup S^3)(\mathbb{S}_0) \to Y$, which is the identity outside of a neighborhood of $\{p_1\} \times D^3$. In the case of the closed three manifold invariants, we had that if $w \in S^3$ was a basepoint, then

$$\phi_* \circ F_{Y \sqcup S^3, \mathbb{S}_0, \overline{\mathfrak{s}}} \circ F_{Y,0} = S^+_{\phi(w)}$$

and a similar relation held for the 4-handle maps and S^-_w. As an analog for links, we prove the following:

Lemma 12.3. *Suppose that $\mathbb{U}$ is an unknot in S^3 and D is a disk bounding $\mathbb{U}$. If ϕ is the diffeomorphism described in the previous paragraph, and D is a disk in S^3 for $\mathbb{U}$, we have the following:*

$$F_{Y \sqcup S^3, \mathbb{S}_0, \overline{\mathfrak{s}}} \circ F_{Y,0,\mathbb{U}} \simeq S^+_{\phi(\mathbb{U}), \phi(D)}.$$

If $\mathbb{S}_2$ denotes the canonical framed 2-sphere in $(Y \sqcup S^3)(\mathbb{S}_0)$, then

$$F_{Y \sqcup S^3, 4, \mathbb{U}} \circ F_{(Y \sqcup S^3)(\mathbb{S}_0), \mathbb{S}_2, \overline{\mathfrak{s}}} \simeq S^-_{\phi(\mathbb{U}), \phi(D)}.$$

Proof. The proof is an easy model computation, nearly identical to the analogous statement for the 3- and 4-manifold invariants (cf. Lemma [Zem15b, Lemma 13.1]). □

To introduce maps for link cobordisms where a component of Σ doesn't have an end in one of the $\mathbb{L}_i$, we pick standard 4-balls along arcs of $\mathcal{A}$. The configuration is shown in Figure 12.1. After removing enough 4-balls, we can use Theorem A′ to construct cobordism maps for the punctured link cobordism, which are invariants (for a given choice of 4-balls).

We need to show that the induced map is independent of the choice of standard 4-balls. It is not the case that any two configurations of 4-balls induce diffeomorphic cobordisms, so we cannot just apply the invariance result of Theorem A′. Instead we show that in general the maps are invariant under removing additional 4-balls along arcs of $\mathcal{A}$. By a sequence of removing additional 4-balls or filling in 4-balls, we can go from one collection to another. To this end, we prove the following lemma:

Lemma 12.4. *Suppose $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a decorated link cobordism with $F = (\Sigma, \mathcal{A})$, and B is a 4-ball intersecting Σ in a disk along an arc of $\mathcal{A}$. Let*

$$(W_0, F_0) : (Y_1 \sqcup S^3, \mathbb{L}_1 \sqcup \mathbb{U}) \to (Y_2, \mathbb{L}_2)$$

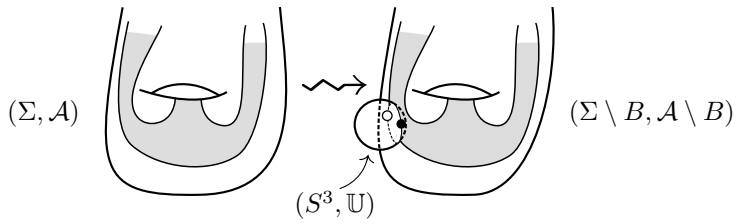

FIGURE 12.1. The process of removing a standard 4-ball from $(W, \Sigma, \mathcal{A})$ which intersect Σ along an arc of $\mathcal{A}$.

be the link cobordism obtained by removing B from W and Σ, then

$$F_{W,F,\mathfrak{s}} \simeq F_{W_0,F_0,\mathfrak{s}_0} \circ F_{Y_1,0,\mathbb{U}}.$$

Proof. This is a model computation. To a Morse function inducing a parametrized Kirby decomposition of (W, F), we introduce a canceling index 0/1 pair such that the index 0 critical point occurs inside of B, along an arc of $\mathcal{A}$, and the index 1 critical point occurs on Σ, in a $\mathbf{z}$-region. By Lemma 12.3, the composition of the 0-handle and the 1-handle map is $S^+_{(U,w,z),D)}$, for an unknot U bounding a disk D in the complement of the rest of the link.

Using the schematic in Figure 12.2, it thus sufficient to show that

$$\phi_* T^-_{w,z} F^{\mathbf{z}}_B S^+_{(U,w,z),D} \simeq \mathrm{id}$$

for the diffeomorphism $\phi : (Y, (L \cup U)(B)) \to (Y, L)$ which is the identity outside of a neighborhood of $B \cup D$. This relation follows from Proposition 7.4 and Lemma 3.13.

□

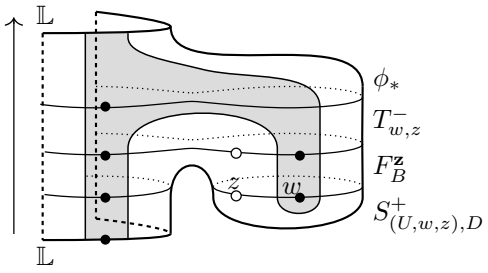

FIGURE 12.2. A schematic showing the composition we need to show is the identity in Lemma 12.4. The dashed lines indicate level sets of a Morse function on the surface.

Of course the dual statement is also true, that if we remove a ball from W along an arc of $\mathcal{A}$ in Σ, and then add the new boundary sphere to the outgoing end $(Y_2, \mathbb{L}_2)$, the map is unaffected:

Lemma 12.5. *If $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is a link cobordism with $F = (\Sigma, \mathcal{A}, \sigma)$, and B is a 4-ball intersecting Σ in a disk along an arc of $\mathcal{A}$, and*

$$(W_0, F_0) : (Y_1, \mathbb{L}_1) \to (Y_2 \sqcup S^3, \mathbb{L}_2 \sqcup \mathbb{U})$$

is the link cobordism obtained by removing B from W and Σ, then

$$F_{W,F,\mathfrak{s}} \simeq F_{Y_2 \sqcup S^3,4,\mathbb{U}} \circ F_{W_0,F_0,\mathfrak{s}_0}.$$

The proof of the above lemma is essentially identical to the proof of Lemma 12.4.

Indeed following this strategy, we can construct maps for arbitrary cobordisms, even where the 4-manifold has empty ends, just as in [Zem15b, Section 13.2]. Summarizing, we can conclude the proof of Theorem A:

Proof of Theorem A. Using Theorem A′, to a link cobordism where not every component of Σ has an incoming or outgoing end, we remove some number of standard 4-balls which intersect Σ along an arc of $\mathcal{A}$, so that each component of Σ has an incoming and outgoing end. To see that the resulting maps are independent of the choice of 4-balls, we use Lemmas 12.4 and 12.5 to see that we can remove as many standard 4-balls as we'd like, along the arcs $\mathcal{A}$, without changing the maps. Through a sequence of removing standard 4-balls along $\mathcal{A}$, or filling them back in, we can go between any two configurations of standard 4-balls. □

13. The composition law

In this Section we prove Theorem B, the Spin^c composition law. Given a link cobordism which is the composition of two link cobordisms, we will take two parametrized Kirby decompositions (one for each link cobordism), and then commute pieces past each other to get a decomposition into elementary parametrized link cobordisms of the composed cobordism, which is almost a parametrized Kirby decomposition, except there are two elementary cobordisms which have 1-dimensional framed links. Most of the Floer homology has already been done, but we have a few lemmas which we need to prove.

Lemma 13.1. *Suppose $\mathbb{L} \subseteq Y$ is a link, and let $\mathbb{L}'$ be the link obtained by adding adjacent basepoints (w, z). If $\mathbb{S} \subseteq Y \setminus L$ is a framed 1-dimensional link, and $\mathfrak{s} \in W(Y, \mathbb{S})$ is a* Spin^c *structure, then*

$$S^{\pm}_{w,z} F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F_{Y,\mathbb{L}',\mathbb{S}_1,\mathfrak{s}} S^{\pm}_{w,z}, \qquad \text{and} \qquad T^{\pm}_{w,z} F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F_{Y,\mathbb{L}',\mathbb{S}_1,\mathfrak{s}} T^{\pm}_{w,z}$$

Proof. This follows by writing $F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$ as $F^{\beta}_{Y,\mathbb{S}_1,\mathfrak{s},\mathcal{B}^\beta}$ for a β-bouquet $\mathcal{B}^\beta$ of the link $\mathbb{S}_1$, and then using [MO10, Prop. 5.2], which is restated using our matrix notation in Lemma 5.1. Of course if one wants to use an α-bouquet $\mathcal{B}^\alpha$ for $\mathbb{S}_1$, then one can use the triangle map computation in Theorem 5.3 instead. □

Lemma 13.2. *Suppose that B is a band for the link $\mathbb{L}$ in Y, and $\mathbb{S}_1$ is a framed 1-dimensional link in Y, which is disjoint from $L \cup B$ and $\mathfrak{s} \in \text{Spin}^c(W(Y, \mathbb{S}))$. If $\mathbf{o} \in \{\mathbf{w}, \mathbf{z}\}$, then*

$$F_{Y,\mathbb{L}(B),\mathbb{S}_1,\mathfrak{s}} F^{\mathbf{o}}_B \simeq F^{\mathbf{o}}_B F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}.$$

Proof. Without loss of generality, suppose that B is an α-band. We now pick a β-bouquet, $\mathcal{B}$, for $\mathbb{S}_1$ in the complement of $B \cup L$. We also pick a Heegaard quadruple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ so that

(1) $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is Heegaard diagram for $(Y, \mathbb{L})$;
(2) $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w}, \mathbf{z})$ is a triple subordinate to the band $B \subseteq Y$;
(3) $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ is a triple subordinate to the bouquet $\mathcal{B}$ in $(Y, \mathbb{L})$;
(4) $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ is a triple subordinate to the band B in $Y(\mathbb{S}_1)$;
(5) $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ is a triple subordinate to the bouquet $\mathcal{B}$ in $(Y, \mathbb{L}(B))$.

Let $\Theta^{\mathbf{o}}_{\alpha'\alpha} \in \mathcal{H}FL^-_{\max,\mathbf{o}}(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha})$ denote the generator. Similarly let $\Theta_{\beta\beta'}$ be the generator of $\mathcal{H}FL^-(\Sigma, \boldsymbol{\beta}, \boldsymbol{\beta}')$ (see Lemma 3.1). The associativity relations applied to the quadruple $(\Sigma, \boldsymbol{\alpha}', \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\beta}', \mathbf{w}, \mathbf{z})$ yield the equivalence

$$F_{\alpha'\alpha\beta',\mathfrak{s}|_{Y(\mathbb{S})}}(\Theta^{\mathbf{o}}_{\alpha'\alpha}, F_{\alpha\beta\beta',\mathfrak{s}}(\mathbf{x}, \Theta_{\beta\beta'})) \simeq F_{\alpha'\beta\beta,\mathfrak{s}}(F_{\alpha'\alpha\beta,\mathfrak{s}|_Y}(\Theta^{\mathbf{o}}_{\alpha'\alpha}, \mathbf{x}), \Theta_{\beta\beta'}),$$

showing that

$$F^{\beta}_{Y,\mathbb{L}(B),\mathbb{S}_1,\mathfrak{s}} F^{\mathbf{o}}_B \simeq F^{\mathbf{o}}_B F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}.$$

Noting that $F_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F^{\alpha}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}} \simeq F^{\beta}_{Y,\mathbb{L},\mathbb{S}_1,\mathfrak{s}}$, by Lemma 4.8, the proof is complete. □

We now prove a slightly weaker version of Theorem B:

Lemma 13.3. *Suppose that $(W, F) : (Y_1, \mathbb{L}_1) \to (Y_2, \mathbb{L}_2)$ is an arbitrary link cobordism and*

$$(W, F) = (W_2, F_2) \circ (W_1, F_1)$$

where each (W_i, F_i) is a link cobordisms such that each component of W has a nonempty incoming and outgoing end, and each component of Σ contains a nonempty incoming and outgoing end. If $\mathfrak{s}_1$ and $\mathfrak{s}_2$ are Spin^c *structures on W_1 and W_2 respectively, then we have*

$$F_{W_2,F_2,\mathfrak{s}_2} \circ F_{W_1,F_1,\mathfrak{s}_1} \simeq \sum_{\substack{\mathfrak{s} \in \text{Spin}^c(W) \\ \mathfrak{s}|_{W_i} = \mathfrak{s}_i}} F_{W,F,\mathfrak{s}}.$$

Proof. Take parametrized Kirby decompositions $\mathcal{K}_1 = \mathcal{K}(f_1, \mathbf{b}_1)$ of (W_1, F_1) and $\mathcal{K}_2 = \mathcal{K}(f_2, \mathbf{b}_2)$ of (W_2, F_2), induced by very nice Morse functions f_1 and f_2 with collections of regular values $\mathbf{b}_1$ and $\mathbf{b}_2$. Also pick gradient like vector fields v_1 and v_2, inducing the decompositions. We can assume that f_1 and f_2 patch together, and v_1 and v_2 patch together to give Morse functions and gradient like flowlines on (W, F). The collections of regular values $\mathbf{b}_1$ and $\mathbf{b}_2$ induce a collection $\mathbf{b}$ for f. This yields a decomposition, $\mathcal{K}(f, \mathbf{b})$, of (W, F) into elementary parametrized cobordisms, which is the one obtained by just stacking $\mathcal{K}_1$ and $\mathcal{K}_2$. We note that $\mathcal{K}(f, \mathbf{b})$ may not be a parametrized Kirby decomposition since the terms are in the wrong order, and there may be two 2-handle cobordisms. Nonetheless, we can still take the composition of the cobordism maps for each elementary parametrized cobordism in the decomposition, which results in just $F_{W_2,F_2,\mathfrak{s}_2} \circ F_{W_1,F_1\mathfrak{s}_1}$.

Notice that the index 1, 2 and 3 dimensional critical points of f_1 which are away from Σ_1 have unstable manifolds which do not intersect Σ, and similarly the critical points of f_2 away from Σ have stable manifolds which don't intersect Σ. Hence, assuming that (f, v) is generic, we can modify the Morse functions to pull all of the critical points of $f_1|_{W_1\setminus\Sigma_1}$ above any critical point of $f_2|_{\Sigma_2}$ along Σ, and then pull the index 3 critical points of f_1 above the index 1 and 2 critical points of f_2, and pull the index 2 critical points of f_1 above the index 1 critical points of f_1.

The composition of the map for each piece of the stacked decomposition is invariant under pulling the 3-handles of f_1 above the 1-handles of f_2 by Lemma 4.4. The composition is also invariant of pulling the 3-handles of f_1 past the 2-handles of f_2 by the triangle map computation of [Zem15b, Theorem 8.8].

Hence it remains to show that the composition is unchanged when we pull the 1-,2- and 3-handles of $f_1|_{W_1\setminus\Sigma_1}$ past the elementary cobordisms associated to critical points of $f_2|_\Sigma$ and $f_2|_\mathcal{A}$. The critical points of $f_2|_\mathcal{A}$ induce maps which are a composition of the basepoint moving maps along the link components, and the quasi-stabilization maps. The maps for surgery on framed spheres away from the link commute with the basepoint moving maps by diffeomorphism invariance of the maps. The 1- and 3-handles of f_1 can be pulled past the critical points of $f_2|_\mathcal{A}$ without changing the composition of the maps by Lemma 4.3 (showing that the 1-handle and 3-handle maps commute with quasi-stabilization). The 1- and 3-handle maps can be commuted past critical points of $f_2|_\Sigma$ using Lemma 4.4 (to commute a 1- or 3-handle past the 1-handle map from the critical point along Σ_2) as well as the triangle map computation from [Zem15b, Theorem 8.8] (to commute a 1- or 3-handle past the band map for the critical point along Σ_2). The 2-handles of f_1 can be commuted past the critical points of $f_2|_\mathcal{A}$ by Lemma 13.1. Similarly the 2-handles of f_1 can be commuted past the critical points of $f_2|_\Sigma$ by Lemma 13.2 (to commute the 2-handle maps past the band maps), and [Zem15b, Theorem 8.8] (to commute the 2-handle maps past the 1-handle map).

Thus we get that $F_{W_2,F_2,\mathfrak{s}_2} \circ F_{W_1,F_1,\mathfrak{s}_1}$ is equal to the composition of maps induced by a decomposition of (W, F) into parametrized elementary cobordisms, which has elementary cobordisms satisfying the requirements of a parametrized Kirby decomposition, except that there are exactly two adjacent terms corresponding to surgery on framed 1-dimensional links. As W is obtained from the union of the two 2-handle pieces by adding 1-handles and 3-handles, a Spin^c structure defined on the union of the two 2-handle cobordisms extends uniquely to a Spin^c structure on all of W. Hence Lemma 4.9 (the composition law for the 2-handle maps) implies the Lemma. □

Using the previous Lemma, we can now prove Theorem B:

Proof Theorem B. This now follows from Lemma 13.3, and the construction of the general cobordism maps for cobordisms with empty ends defined in Section 12 using the same strategy as the final component of the proof of [Zem15b, Theorem E] found in Subsection 13.3 of that paper. See also [Juh15a, Theorem 11.3]. □

14. Link functors and graph TQFTs

In this section we prove Theorem C by investigating the relationship between the graph TQFT defined in [Zem15b] and the maps defined in this paper. We will first develop some notation, and then restate Theorem C more precisely using this notation. In the following subsections we will prove Theorem C.

In [Zem15a] and[Zem15b], the author develops cobordism maps in Heegaard Floer homology for 4-manifolds with embedded graphs. In this section we consider the relationship between the graph cobordism maps developed there, and the link cobordism maps defined in this paper. The graph cobordism maps

defined in [Zem15b] are defined using relative homology maps A_λ, which count Maslov index one disks ϕ with a multiplicity depending on how the $\boldsymbol{\alpha}$-boundary of the homology class ϕ interacts with the curve λ, which we view as embedded in the Heegaard surface Σ. We denote the cobordism maps defined in [Zem15b] by

$$F^A_{W,\Gamma,\mathfrak{s}},$$

to emphasize that they were defined with the maps A_λ. As a variation, one can consider maps

$$F^B_{W,\Gamma,\mathfrak{s}},$$

defined using the relative homology maps B_λ, instead of A_λ. In [HMZ], we show that if the vertices of Γ have valence at most 3, then $F^B_{W,\Gamma,\mathfrak{s}} = F^A_{W,\Gamma^{\text{opp}},\mathfrak{s}}$, where Γ^{opp} has the same underlying graph as Γ, but has all cyclic orders reversed.

We now describe a surface with divides, $F(\Gamma) \subseteq W$, associated to a cyclically ordered graph in Γ. First, give each edge of Γ an orientation (arbitrarily). If p is a vertex in the interior of W, first isotope Γ near p so that the edges adjacent p are locally embedded in an oriented disk D^2 centered at p compatibly with the cyclic ordering. Pick a normal vector field v_1 on D, viewing D as a submanifold of W. If e is an edge of Γ, we can assign $e \cap D$ a normal vector $v_2 \in TD|_{e\cap D}$, viewing $e \cap D$ as an oriented submanifold of D. This determines a 2-frame of vector fields, (v_1, v_2), on each $e \cap D$, with both v_1 and v_2 being orthogonal to e.

We now wish to extend the frames (v_1, v_2) over all edges, so that the along an edge, the vector fields v_1 and v_2 are orthogonal to the tangent space of the edge. The obstruction to extending the 2-frame along an edge lies in $\pi_2(V(3,2))$, where $V(3,2)$ is the Stiefel manifold of 2-frames in $\mathbb{R}^3$. Hence the ambiguity in defining this frame lies in $\pi_1(V(3,2)) \cong \mathbb{Z}_2$, so the 2-frames aren't determined uniquely, even up to isotopy. However this ambiguity is minor, and we will give a brief discussion in Subsection 14.1.

To define a surface with divides $F(\Gamma)$, one starts with the disks D centered at each vertex. In the direction of v_2, one thickens the edges e to strips, which near the vertices are embedded in the disks. The union of all of the disks and strips is an embedded surface with boundary $\Sigma_0(\Gamma)$. One then pushes off two copies of $\Sigma_0(\Gamma)$ in the $+v_1$ and $-v_1$ directions, while remaining fixed on $\partial\Sigma_0(\Gamma)$. This yields two surfaces, $\Sigma_0^+(\Gamma)$ and $\Sigma_0^-(\Gamma)$, with common boundary along the boundary of $\Sigma_0(\Gamma) \cap \operatorname{int} W$. This intersection is naturally a collection of arcs, $\mathcal{A}$, which form a dividing set. We give $\Sigma_0^+(\Gamma)$ the designation of being type $\mathbf{w}$, and $\Sigma_0^-(\Gamma)$ the designation of being type $\mathbf{z}$. The orientable surface $\Sigma(\Gamma)$ is given the orientation that agrees with Σ_0^+.

Define $U_1 = \Sigma(\Gamma) \cap Y_1$ and $U_2 = \Sigma(\Gamma) \cap Y_2$. We set

$$\mathcal{A}(\Gamma) = (\partial\Sigma_0^+(\Gamma)) \cap \operatorname{int} W = (\partial\Sigma_0^-(\Gamma)) \cap \operatorname{int} W.$$

We define

$$F(\Gamma) = (\Sigma(\Gamma), \mathcal{A}(\Gamma)),$$

a surface with divides. We give $\partial\Sigma(\Gamma)$, a collection of unknots, the orientation induced by an outward vector field (as is our standard convention).

Given a coloring $(\sigma, \mathfrak{P})$ of Γ, we define a coloring $(\sigma, \mathfrak{P}_W \sqcup \mathfrak{P}_Z)$ on the surface with divides $F(\Gamma)$, simply by $\mathfrak{P}_W = \mathfrak{P} \times \{W\}$ and $\mathfrak{P}_Z = \mathfrak{P} \times \{Z\}$ (here W and Z are just letters). We color a component of $\Sigma_0(\Gamma)^+$ to the corresponding color of the graph component in $\mathfrak{P}_W$, and color components of $\Sigma_0(\Gamma)^-$ using the corresponding color in $\mathfrak{P}_Z$. We put a single $\mathbf{w}$ and $\mathbf{z}$ basepoint in each component of U_i, the $\mathbf{w}$-basepoint in $\Sigma_0(\Gamma)^+$ and the $\mathbf{z}$-basepoint in $\Sigma_0(\Gamma)^-$.

We define a colored surface with divides $\overline{F}(\Gamma)$ to have the same underlying surface with divides as $F(\Gamma)$, but with the roles of $\Sigma_0(\Gamma)^+$ and $\Sigma_0(\Gamma)^-$ reversed (i.e. designate $\Sigma_0(\Gamma)^+$ as type-$\mathbf{z}$, and $\Sigma_0(\Gamma)^-$ as type-$\mathbf{w}$).

Finally note that if $(Y_1, \mathbf{w}_1)$ and $(Y_2, \mathbf{w}_2)$ denote the ends of (W, Γ), the isotopies of $\Sigma_0(\Gamma)$ to $\Sigma_0^\pm(\Gamma)$ span disks in Y_1 and Y_2, which are centered at the points $\mathbf{w}_1$ and $\mathbf{w}_2$. The disks induce a diffeomorphism ϕ_i, well defined up to isotopy and equal to the identity outside a neighborhood of these disks, corresponding to moving the basepoints $\mathbf{w}_i$ to the $\mathbf{w}$-basepoints of U_i for the surface $F(\Gamma)$. Similarly there are analogous diffeomorphisms $\overline{\phi}_i$ corresponding to moving the $\mathbf{w}_i$ basepoints to the $\mathbf{w}$-basepoints on U_i for the surface $\overline{F}(\Gamma)$.

We have the following:

Theorem C. *Suppose that (W,Γ) is a graph cobordism and each vertex of Γ has valence at most three. Suppose that $F(\Gamma)$ is a colored surface with divides, constructed as above. After setting all of the $\mathfrak{P}_Z$ variables to 1, we have an equalities*

$$F_{W,F(\Gamma),\mathfrak{s}} = (\phi_2)_* F^B_{W,\Gamma,\mathfrak{s}}(\phi_1^{-1})_* \qquad \text{and} \qquad F_{W,\overline{F}(\Gamma),\mathfrak{s}} = (\overline{\phi}_2)_* F^A_{W,\Gamma,\mathfrak{s}}(\overline{\phi}_1^{-1})_*.$$

The claim is expected to hold in the case that Γ has vertices of valence more than three, though we will not prove it.

14.1. **The ambiguity of $(\Sigma(\Gamma),\mathcal{A}(\Gamma))$.** As mentioned before, the surfaces with divides $(\Sigma(\Gamma),\mathcal{A}(\Gamma))$ are not quite specified up to isotopy. Indeed as we mentioned before there is an ambiguity corresponding to twisting along the edges of the graph with an element of $\pi_1(V(3,2))$. A homotopy of the oriented 2-frame (v_1,v_2) through 2-frames used in the construction yields an isotopy of the resulting surface with divides. Given two choices of 2-frames, we can always homotope one so it agrees with the other over the disks. Along the edges, the resulting surface $\Sigma(\Gamma)$ depends only oriented subspaces of $(Te)^{\perp}$ spanned by the 2-frame, whereas the divides depend on the choice of 2-frames along the edges. Noting that the oriented Grassmannian is homeomorphic to S^2, and is simply connected, we conclude that the underlying surfaces have no ambiguity, but the decorated surface has the $\mathbb{Z}_2$ ambiguity along each edge already described.

The generator of $\pi_1(V(3,2))$ is easy to describe, it is the result of taking a fixed frame, and having $SO(2)\cong S^1$ act on that frame, by rotations. In particular, the ambiguity of the decorated surface $(\Sigma(\Gamma),\mathcal{A}(\Gamma))$ is only through twisting the divides along the surface, along an edge.

The effect to the link cobordism maps is to insert the diffeomorphism map corresponding to a twist around a link component, at various stages. We can use Sarkar's formula from [Sar11a] (proven in [Zem16] for the versions of link Floer homology we are working with) to see that this action is trivial, though it is easier to directly observe that the action is trivial for an unknotted component with exactly two basepoints since we can pick a Heegaard diagram where the w and z basepoints are adjacent to each other, and not separated by any $\boldsymbol{\alpha}$- or $\boldsymbol{\beta}$- curves. For such a diagram the diffeomorphism action corresponding to Sarkar's action is obviously trivial.

Another observation is that Sarkar's map on link Floer homology is always an involution for algebraic reasons (since Φ_w^2 and Ψ_z^2 are chain homotopic to zero). By our observations about $\pi_1(V(3,2))$, the corresponding decorated knot cobordism for the unknot is also an involution as a decorated knot cobordism. It would be interesting to know if there are other knots for which the knot cobordism is an involution, or if there were other topological reasons for this map to be an involution.

14.2. **The maps Ψ_z, A_λ and the proof of Theorem C.** Our goal in this section is to prove a relation between the relative homology maps A_λ, defined in [Zem15b] on $CF^\circ(Y,\mathbf{w},\mathfrak{s})$ and the maps Ψ_z defined in this paper on curved link Floer homology. We recall that for an immersed curve λ in Σ from w_1 to w_2, we defined maps

$$A_\lambda : CF^\circ(Y,\mathbf{w},\mathfrak{s}) \to CF^\circ(Y,\mathbf{w},\mathfrak{s})$$

which we defined by the formula

$$A_\lambda(\mathbf{x}) = \sum_{\mathbf{y}\in\mathbb{T}_\alpha\cap\mathbb{T}_\beta} \sum_{\substack{\phi\in\pi_2(\mathbf{x},\mathbf{y})\\ \mu(\phi)=1}} a(\lambda,\phi)\#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)}\cdot\mathbf{y},$$

where $a(\lambda,\phi)$ is the sum of all of the changes of multiplicities of ϕ across the $\boldsymbol{\alpha}$-curves as one traverses along λ. The map A_λ satisfies

$$\partial A_\lambda + A_\lambda\partial = U_{w_1} + U_{w_2}.$$

When we give w_1 and w_2 the same color, these become chain maps. The maps A_λ were used in the construction of the graph cobordism maps in [Zem15b].

Lemma 14.1. *Suppose that w_1, z, and w_2 are consecutive basepoints on a link $\mathbb{L}$ (in that order, though w_1 and w_2 need not be distinct), and λ is the subarc of L from w_1 to w_2. Then after identifying all of the $V_{\mathbf{z}}$-variables with 1, we have*

$$\Psi_z \simeq A_\lambda + U_{w_2}\Phi_{w_2} \simeq B_\lambda + U_{w_1}\Phi_{w_1}.$$

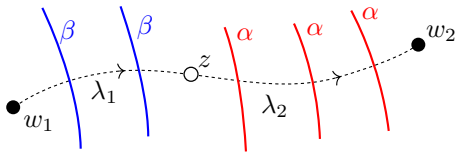

FIGURE 14.1. The curve λ is the concatenation of λ_1 and λ_2.

Proof. Write λ as the concatenation of two paths, λ_1 and λ_2, as in Figure 14.1.

As λ is a subarc of L, we can assume that λ_1 does not intersect any of the $\boldsymbol{\alpha}$-curves, and that λ_2 does not intersect any of the $\boldsymbol{\beta}$-curves. Hence when adding up the changes across the $\boldsymbol{\alpha}$-curves to compute $a(\lambda, \phi)$, we will count only the differences of intersections of λ_2 with the $\boldsymbol{\alpha}$-curves. Since λ_2 also does not intersect any $\boldsymbol{\beta}$-curves, the quantity $a(\lambda, \phi)$ will just be the total change of the multiplicities of ϕ along λ_2, i.e.,

$$a(\lambda, \phi) = n_{w_2}(\phi) - n_z(\phi).$$

Multiplying $\#\widehat{\mathcal{M}}(\phi) U_{\mathbf{w}}^{n_{\mathbf{w}}(\phi)} \cdot \mathbf{y}$ by this quantity and summing over all disks, we get that

$$A_\lambda(\mathbf{x}) = U_{w_2} \Phi_{w_2}(\mathbf{x}) + \Psi_z(\mathbf{x}),$$

for this diagram. Analogously, the map B_λ computes differences across the $\boldsymbol{\beta}$-curves, and the same line of reasoning yields

$$B_\lambda(\mathbf{x}) = U_{w_1} \Phi_{w_1}(\mathbf{x}) + \Psi_z(\mathbf{x}).$$

Note that these two formulas are consistent with the relation $A_\lambda + B_\lambda = U_{w_1}\Phi_{w_1} + U_{w_2}\Phi_{w_2}$. □

Proof of Theorem C. The maps obviously agree for path cobordisms, since both are defined only using the handle attachment maps. Similarly the 0- and 4-handle maps agree, once we identify all $\mathfrak{P}_Z$ variables with 1. Hence using the composition law, it is sufficient to show the claim for graphs in cylinders of the form $Y \times [0,1]$. By the composition law for both types of cobordism maps, it is sufficient to show the claim for elementary graphs with valence at most 3. For elementary graphs with a valence 1 vertex, the claim follows since the graph action map is $S^{\pm}_w$ and the surface cobordism map is the disk-stabilization map $S^{\pm}_{(U,w,z),D}$, which becomes $S^{\pm}_w$ after setting the $V_{\mathbf{z}}$-variables all to 1. For a graph with a single interior vertex of valence 2, connected to vertices on both $Y \times \{0\}$ and $Y \times \{1\}$, both maps trivially agree since they are both equal to the diffeomorphism map corresponding to moving the basepoint according to the path (note that according to [Zem15b, Lemma 7.8], the graph cobordism maps are invariant under subdividing an edge)

Similarly, elementary graphs with an edge between two incoming vertices, or two outgoing vertices can be turned into a graph with a trivalent vertex, and a valence 1 vertex, by adding a trivial strand (cf. [Zem15b, Lemma 7.11]).

The only type of elementary graph remaining is an elementary flow graph with a single valence-3 vertex. We first consider the case that there are two outgoing vertices, and one incoming vertex. This is shown in Figure 14.2. We will abuse notation, and identity the variables $\mathbf{w}_i$ with the corresponding $\mathbf{w}$-variables in the links (which are obtained by pushing $\mathbf{w}_i$ along a diffeomorphism), and in doing so we will not include the diffeomorphisms ϕ_1 or ϕ_2 in the notation.

We will compute the map, under the identification of $w \in V_{\text{in}}$ with $w \in V_{\text{out}}$ via the diffeomorphism which moves w along the two edges in the graph.

We use [HMZ, Lemma 5.5][7], which shows that

$$F^B_{Y\times[0,1],\Gamma} \simeq B_\lambda S^+_{w'} + U_w S^+_{w'} \Phi_w. \tag{14}$$

On the other hand, using the decomposition of surface in Figure 14.2, we see that the surface map for the surface F, pictured, is

$$F_{Y\times[0,1],F} = F^{\mathbf{z}}_{B'} \Psi_{z'} F^{\mathbf{z}}_B S^+_{(U,w',z'),D},$$

[7]Note that in [HMZ], this formula was stated for $F^A_{Y\times[0,1],\Gamma}$, but the same proof holds for the type B maps.

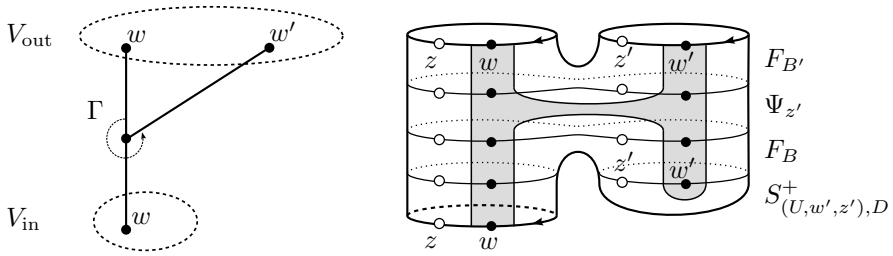

FIGURE 14.2. A cyclically ordered trivalent elementary graph, and the corresponding surface, as well as the decomposition of the surface into pieces. Here we compute the maps induced after identifying $w \in V_{\text{in}}$ with $w \in V_{\text{out}}$ via the diffeomorphism induced by moving w along the two edges between them. Note also the orientations on the boundary components, which are given by our convention for surfaces to be oriented by an outward normal vector.

for bands B and B'. We note that after setting all of the $V_{\mathbf{z}}$-variables equal to 1, the maps $F^{\mathbf{z}}_{B'}$ and $F^{\mathbf{z}}_{B}$ become simply change of diagrams maps. In particular, after setting $V_{\mathbf{z}} = 1$, as maps on the coherent chain homotopy type invariant CF°, using Lemma 14.1, we get

$$F_{Y\times[0,1],F} \simeq F_{B'}\Psi_{z'}F_B S^+_{(U,w',z'),D} \simeq \mathrm{id}\circ(B_\lambda + U_w\Phi_w)\circ \mathrm{id}\circ S^+_{w'} \simeq B_\lambda S^+_{w'} + U_w\Phi_w S^+_{w'},$$

which is equal to our expression for $F^B_{Y\times[0,1],\Gamma}$ in Equation (14).

The computation in the case that there are two incoming vertices and one outgoing vertex is proven analogously, using [HMZ, Lemma 5.4] for the computation of the graph action map.

Finally we consider elementary flow graphs with a valence three vertex such that all three adjacent edges connect to vertices on the incoming end, or all three connect to vertices on the outgoing end. We note that by modifying the Cerf decomposition of the graph, such an elementary flow graph can be split into a trivalent flow graph of the type already considered and a flow graph with an edge connecting two incoming vertices, or two outgoing vertices (cf. [HMZ, Figure 5]). Since we have already considered such elementary flow graphs, the proof is complete. □

REFERENCES

[AE] A. Alishahi and E. Eftekhary, *Tangle Floer homology and cobordisms between tangles*, (preprint).

[Cer70] J. Cerf, *La stratification naturelle des espaces de fonctions diff'erentiables réelles et le théorè de la pseudo-isotopie*, Publications matheématiques de l'I.H.É.S **39** (1970), 5-173.

[GWW] D. Gay, K. Wehrheim, and C. Woodward, *Connected Cerf Theory*, (preprint).

[Hat73] A. Hatcher, *Pseudoisotopies of compact manifolds*, Asterisque **6** (1973).

[HKM08] K. Honda, W. Kazez, and G. Matić, *Contact structures, sutured Floer homology and TQFT* (2008), arXiv:0807.2431 [math.GT].

[HM15] K. Hendricks and C. Manolescu, *Involutive Heegaard Floer homology*, arXiv:1507.00383[math.GT] (2015).

[HMZ] K. Hendricks, C. Manolescu, and I. Zemke, *A connected sum formula for involutive Heegaard Floer homology*, arxiv:1607.07499[math.GT].

[JM15] A. Juhász and M. Marengon, *Concordance maps in knot Floer homology* (2015), arXiv:1509.02738[math.GT].

[JM16] ———, *Computing cobordism maps in link Floer homology and the reduced Khovanov TQFT* (2016), arXiv:1503.00665[math.GT].

[JT12] A. Juhász and D. P. Thurston, *Naturality and mapping class groups in Heegaard Floer homology* (2012), arXiv:1210.4996 [math.GT].

[Juh15a] A. Juhász, *Cobordisms of sutured manifolds and the functoriality of link Floer homology* (2015), arXiv:0910.4382v4[math.GT].

[Juh15b] ———, *Defining and classifying TQFTs via surgery* (2015), arXiv:1408.0668v5[math.GT].

[KR08] M. Khovanov and L. Rozansky, *Matrix factorizations and link homology*, Fund. Math. **199** (2008), 1-91.

[Kir78] R. Kirby, *A calculus for framed links in S^3*, Invent. Math **45** (1978), 35-56.

[Kos93] A. Kosinski, *Differential Manifolds*, Academic Press, San Diego, CA, 1993.

[Lip06] R. Lipshitz, *A cylindrical reformulation of Heegaard Floer homology*, Geom. Topol. **10** (2006), 955-1097.

[Man07] C. Manolescu, *An unoriented skein exact triangle for knot Floer homology*, Math. Res. Let. **14** (2007), 839-852.

[MO10] C. Manolescu and P. Ozsváth, *Heegaard Floer homology and integer surgeries on links* (2010), arXiv:1011.1317v3 [math.GT].
[Nic03] L. Nicolaescu, *The Reidemeister Torsion of 3-manifolds*, De Gruyter studies in Mathematics, New York, NY, 2003.
[OSS15] P. S. Ozsváth, A. Stipsicz, and Z. Szabó, *Unoriented knot Floer homology and the unoriented four-ball genus* (2015), arXiv:1508.03243 [math.GT].
[OS04a] P. S. Ozsváth and Z. Szabó, *Holomorphic disks and topological invariants for closed three manifolds*, Ann. of Math. (2) **159** (2004), no. 3, 1027-1158.
[OS04b] ———, *Holomorphic disks and knot invariants*, Adv. Math **186** (2004), 58-116.
[OS06] ———, *Holomorphic triangles and invariants for smooth four-manifolds*, Adv. in Math. **202,** (2006), Issue 2, 326-400.
[OS08] ———, *Holomorphic disks, link invariants, and the multi-variable Alexander polynomial*, Alg. & Geom. Top. **8** (2008), 615-692.
[OS04] ———, *Holomorphic disks and three-manifold invariants: properties and applications*, Ann. of Math (2) **159** (2004), no. 3, 1159-1245.
[Pal63] R. Palais, *Morse Theory on Hilbert manifolds*, Topology **2** (1963), 299-340.
[Ras03] J. Rasmussen, *Floer homology and knot complements* (2003), arXiv:0306378 [math.GT].
[Sar11a] S. Sarkar, *Moving basepoints and the induced automorphisms of link Floer homology* (2011), arXiv:1109:2168 [math.GT].
[Sar11b] ———, *Grid diagrams and the Ozsváth-Szabó tau-invariant*, Math. Res. Let. **18** (2011), Number 6, 1239-1257.
[Zem15a] I. Zemke, *A graph TQFT for hat Heegaard Floer homology* (2015), arXiv:1010.2808 [math.GT].
[Zem15b] ———, *Graph cobordisms and Heegaard Floer homology* (2015), arXiv:1512.01184 [math.GT].
[Zem16] ———, *Quasi-stabilization and basepoint moving maps in link Floer homology* (2016), arxiv:1604.04316 [math.GT].
[Zema] ———, *Absolute gradings in link Floer homology*, (in preparation).
[Zemb] ———, *A connected sum formula for involutive link Floer homology*, (in preparation).

Department of Mathematics, UCLA, 520 Portola Plaza, Los Angeles, CA 90095, USA
E-mail address: ianzemke@math.ucla.edu